39XX/51XX Service Delivery and Aggregation Switches

Configuration

Release 6.10

What's inside...

Configuration fundamentals
Configuration management
Port management
Hardware resource management
Synchronization configuration
Link Layer Discovery Protocol (LLDP) configuration
VLAN management
IP management
MEF L2 VPN configuration
Provider Backbone Bridge Traffic Engineering (PBB-TE) implementation
Multiprotocol Label Switching (MPLS) Layer 2 VPN configuration
L2 Control Frame Tunneling configuration
Quality of Service configuration
Multicast services configuration
Appendix A: Error codes

009-3203-008 - Standard Revision F
September 2013

Copyright© 2012-2013 Ciena® Corporation. All rights reserved.

LEGAL NOTICES

THIS DOCUMENT CONTAINS CONFIDENTIAL AND TRADE SECRET INFORMATION OF CIENA CORPORATION AND ITS RECEIPT OR POSSESSION DOES NOT CONVEY ANY RIGHTS TO REPRODUCE OR DISCLOSE ITS CONTENTS, OR TO MANUFACTURE, USE, OR SELL ANYTHING THAT IT MAY DESCRIBE. REPRODUCTION, DISCLOSURE, OR USE IN WHOLE OR IN PART WITHOUT THE SPECIFIC WRITTEN AUTHORIZATION OF CIENA CORPORATION IS STRICTLY FORBIDDEN.

EVERY EFFORT HAS BEEN MADE TO ENSURE THAT THE INFORMATION IN THIS DOCUMENT IS COMPLETE AND ACCURATE AT THE TIME OF PUBLISHING; HOWEVER, THE INFORMATION CONTAINED IN THIS DOCUMENT IS SUBJECT TO CHANGE.

While the information in this document is believed to be accurate and reliable, except as otherwise expressly agreed to in writing CIENA PROVIDES THIS DOCUMENT "AS IS" WITHOUT WARRANTY OR CONDITION OF ANY KIND, EITHER EXPRESS OR IMPLIED. The information and/or products described in this document are subject to change without notice.

Copyright© 2012-2013 Ciena® Corporation. All Rights Reserved

The material contained in this document is also protected by copyright laws of the United States of America and other countries. It may not be reproduced or distributed in any form by any means, altered in any fashion, or stored in a data base or retrieval system, without express written permission of the Ciena Corporation.

Security

Ciena® cannot be responsible for unauthorized use of equipment and will not make allowance or credit for unauthorized use or access.

Contacting Ciena

Corporate Headquarters	410-694-5700 or 800-921-1144	www.ciena.com
Customer Technical Support/Warranty		
In North America	1-800-CIENA24 (243-6224) 410-865-4961	
In Europe, Middle East, and Africa	800-CIENA-24-7 (800-2436-2247) +44-207-012-5508	
In Asia-Pacific	800-CIENA-24-7 (800-2436-2247) +81-3-6367-3989 +91-124-4340-600	
In Caribbean and Latin America	800-CIENA-24-7 (800-2436-2247) 410-865-4944 (USA)	
Sales and General Information	410-694-5700	E-mail: sales@ciena.com
In North America	410-694-5700 or 800-207-3714	E-mail: sales@ciena.com
In Europe	+44-207-012-5500 (UK)	E-mail: sales@ciena.com
In Asia	+81-3-3248-4680 (Japan)	E-mail: sales@ciena.com
In India	+91-124-434-0500	E-mail: sales@ciena.com
In Latin America	011-5255-1719-0220 (Mexico City)	E-mail: sales@ciena.com
Training	877-CIENA-TD (243-6283) or 410-865-8996	E-mail: techtng@ciena.com

For additional office locations and phone numbers, please visit the Ciena web site at www.ciena.com.

IMPORTANT: PLEASE READ THIS LICENSE AGREEMENT (“AGREEMENT”) CAREFULLY BEFORE INSTALLING OR USING CIENA CORPORATION (“Ciena”) SOFTWARE, HARDWARE OR DOCUMENTATION (COLLECTIVELY, THE “EQUIPMENT”).

BY INSTALLING OR USING THE EQUIPMENT, YOU ACKNOWLEDGE THAT YOU HAVE READ THIS AGREEMENT AND AGREE TO BE BOUND BY ITS TERMS AND CONDITIONS.

1. Right to Use License; Restrictions. Subject to these terms, and the payment of all applicable license fees, Ciena grants to you, as end user, a non-exclusive license to use the Ciena software (the “Software”) in object code form solely in connection with, and as embedded within, the Equipment,. You shall have the right to use the Software solely for your own internal use and benefit. You may make one copy of the Software and documentation solely for backup and archival purpose, however you must reproduce and affix all copyright and other proprietary rights notices that appear in or on the original. You may not, without Ciena's prior written consent, (i) sublicense, assign, sell, rent, lend, lease, transfer or otherwise distribute the Software; (ii) grant any rights in the Software or documentation not expressly authorized herein; (iii) modify the Software nor provide any third person the means to do the same; (iv) create derivative works, translate, disassemble, recompile, reverse engineer or attempt to obtain the source code of the Software in any way; or (v) alter, destroy, or otherwise remove any proprietary notices or labels on or embedded within the Software or documentation. You acknowledge that this license is subject to Section 365 of the U.S. Bankruptcy Code and requires Ciena's consent to any assignment related to a bankruptcy proceeding. Sole title to the Software and documentation, to any derivative works, and to any associated patents and copyrights, remains with Ciena or its licensors. Ciena reserves to itself and its licensors all rights in the Software and documentation not expressly granted to you. You shall preserve intact any notice of copyright, trademark, logo, legend or other notice of ownership from any original or copies of the Software or documentation.

2. Audit: Upon Ciena's reasonable request, but not more frequently than annually without reasonable cause, you shall permit Ciena to audit the use of the Software at such times as may be mutually agreed upon to ensure compliance with this Agreement.

3. Confidentiality. You agree that you will receive confidential or proprietary information (“Confidential Information”) in connection with the purchase, deployment and use of the Equipment. You will not disclose Confidential Information to any third party without prior written consent of Ciena, will use it only for purposes for which it was disclosed, use your best efforts to prevent and protect the contents of the Software from unauthorized disclosure or use, and must treat it with the same degree of care as you do your own similar information, but with no less than reasonable care. You acknowledge that the design and structure of the Software constitute trade secrets and/or copyrighted materials of Ciena and agree that the Equipment is Confidential Information for purposes of this Agreement.

4. U.S. Government Use. The Software is provided to the Government only with restricted rights and limited rights. Use, duplication, or disclosure by the Government is subject to restrictions set forth in FAR Sections 52-227-14 and 52-227-19 or DFARS Section 52.227-7013(C)(1)(ii), as applicable. The Equipment and any accompanying technical data (collectively “Materials”) are commercial within the meaning of applicable Federal acquisition regulations. These Materials were developed fully at private expense. U.S. Government use of the Materials is restricted by this Agreement, and all other U.S. Government use is prohibited. In accordance with FAR 12.212 and DFAR Supplement 227.7202, software delivered to you is commercial computer software and the use of that software is further restricted by this Agreement.

5. Term of License. This license is effective until terminated. Customer may terminate this license at any time by giving written notice to Ciena [or] and destroying or erasing all copies of Software including any documentation. Ciena may terminate this Agreement and your license to the Software immediately by giving you written notice of termination in the event that either (i) you breach any term or condition of this Agreement or (ii) you are wound up other than voluntarily for the purposes of amalgamation or reorganization, have a receiver appointed or enter into liquidation or bankruptcy or analogous process in your home country. Termination shall be without prejudice to any other rights or remedies Ciena may have. In the event of any termination you will have no right to keep or use the Software or any copy of the Software for any purpose and you shall destroy and erase all copies of such Software in its possession or control, and forward written certification to Ciena that all such copies of Software have been destroyed or erased.

6. Compliance with laws. You agree to comply with all applicable laws, including all import regulations, and to obtain all required licenses and permits related to installation and use of Equipment. Software, including technical data, is subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Customer agrees to comply strictly with all such regulations and acknowledges that it has the responsibility to obtain licenses to export, re-export, or import Software.

Copyright© 2012-2013 Ciena® Corporation

7. Limitation of Liability. ANY LIABILITY OF Ciena SHALL BE LIMITED IN THE AGGREGATE TO THE AMOUNTS PAID BY YOU FOR THE SOFTWARE. THIS LIMITATION APPLIES TO ALL CAUSES OF ACTION, INCLUDING WITHOUT LIMITATION BREACH OF CONTRACT, BREACH OF WARRANTY, NEGLIGENCE, STRICT LIABILITY, MISREPRESENTATION AND OTHER TORTS. THE LIMITATIONS OF LIABILITY DESCRIBED IN THIS SECTION ALSO APPLY TO ANY THIRD-PARTY SUPPLIER OF Ciena. NEITHER Ciena NOR ANY OF ITS THIRD-PARTY SUPPLIERS SHALL BE LIABLE FOR ANY INJURY, LOSS OR DAMAGE, WHETHER INDIRECT, SPECIAL, INCIDENTAL OR CONSEQUENTIAL INCLUDING WITHOUT LIMITATION ANY LOST PROFITS, CONTRACTS, DATA OR PROGRAMS, AND THE COST OF RECOVERING SUCH DATA OR PROGRAMS, EVEN IF INFORMED OF THE POSSIBILITY OF SUCH DAMAGES IN ADVANCE

8. General. Ciena may assign this Agreement to any Ciena affiliate or to a purchaser of the intellectual property rights in the Software, but otherwise neither this Agreement nor any rights hereunder may be assigned nor duties delegated by either party, and any attempt to do so will be void. This Agreement shall be governed by the laws of the State of Maryland (without regard to the conflict of laws provisions) and shall be enforceable in the courts of Maryland. The U.N. Convention on Contracts for the International Sale of Goods shall not apply hereto. This Agreement constitutes the complete and exclusive statement of agreement between the parties relating to the license for the Software and supersedes all proposals, communications, purchase orders, and prior agreements, verbal or written, between the parties. If any portion hereof is found to be void or unenforceable, the remaining provisions shall remain in full force and effect.

Copyright© 2012-2013 Ciena® Corporation

Publication history

August 2012

Standard Revision A

First Standard release of document.

December 2012

Standard Revision B

Document updated for Release 6.10.1.

February 2013

Standard Revision C

Document updated for Release 6.10.2.

March 2013

Standard Revision D

Changes are:

- Updated Caution statement in “Restoring default configurations” on page 2-7.
- Updated parameter “priority” to “override-priority” in “Synchronization configuration” on page 5-1.
- Added figure “Policy set structure” on page 9-5

April 2013

Standard Revision E

Changes are:

- 3930 External Timing platform to 3930 Sync + External Timing platform
- 5150 with 5102 module platform to 5150 Standard with BITS platform
- 5102 module to BITS module

Copyright© 2012-2013 Ciena® Corporation

September 2013

Standard Revision F

Updated descriptions of the untagged-data-vid attribute to specify that the untagged-data-vid attribute pushes the customer VID on ingress but does not pop the customer VID on egress on the 3902, 3911, 3920, 3940 and 5140 platforms.

Copyright© 2012-2013 Ciena® Corporation

Contents

About this document xvii

Configuration fundamentals 1-1

Ports 1-1
Hardware resources 1-1
Synchronization configuration 1-2
Link Layer Discovery Protocol 1-2
Virtual Local Area Networks 1-2
IP management 1-3
Metro Ethernet Forum Layer 2 Virtual Private Networks 1-3
Provider Backbone Bridge Traffic Engineering 1-4
Multiprotocol Label Switching 1-4
Layer 2 control frame tunneling 1-4
Quality of Service 1-5
Multicast services 1-6

Configuration management 2-1

Accessing the CLI 2-1
Configuration files 2-1

List of procedures

2-1 Saving configuration changes 2-3
2-2 Displaying configuration files 2-4
2-3 Augmenting the current configuration 2-5
2-4 Restoring default configurations 2-7
2-5 Setting the default configuration files 2-8
2-6 Displaying the default configuration files 2-9
2-7 Resetting default configuration files to factory default files 2-10

Port management 3-1

Port attributes 3-1
Port statistics 3-5
Transceivers 3-8
  Identification 3-8
  Diagnostics 3-9

List of procedures

3-1 Setting port attributes 3-12

Copyright© 2012-2013 Ciena® Corporation

3-2 Resetting port attributes to default 3-13
3-3 Disabling a port 3-14
3-4 Enabling a port 3-15
3-5 Displaying port attributes 3-16
3-6 Displaying port statistics 3-17
3-7 Monitoring port statistics 3-22
3-8 Clearing current statistics 3-27
3-9 Displaying blade information 3-28
3-10 Displaying port capabilities 3-32
3-11 Displaying port Ethernet configuration 3-34
3-12 Displaying port status 3-35
3-13 Displaying a list of supported optics 3-36
3-14 Displaying transceiver information 3-37
3-15 Determining transceiver speed 3-40
3-16 Tuning XFP transceivers 3-42
3-17 Setting the port connector mode 3-44

Hardware resource management 4-1

List of procedures
4-1 Configuring resources 4-7
4-2 Freeing all accelerated CFM over PBB-TE resources 4-9
4-3 Restoring accelerated CFM over PBB-TE resources to default values 4-11
4-4 Freeing all broadcast containment resources 4-13
4-5 Restoring broadcast containment resources to default values 4-14
4-6 Freeing CFM resources 4-16
4-7 Restoring CFM resources to default values 4-17
4-8 Freeing DHCP relay resources 4-19
4-9 Restoring DHCP relay resources to default values 4-20
4-10 Freeing traffic profiling resources 4-22
4-11 Setting traffic profiling resources 4-24
4-12 Restoring traffic profiling resources to default values 4-27
4-13 Freeing virtual circuit statistics resources 4-29
4-14 Restoring virtual circuit statistics resources to default values 4-30
4-15 Configuring virtual switch L2 enhanced transform resources 4-32
4-16 Freeing all virtual switch L2 enhanced transform resources 4-34
4-17 Configuring transport OAM resources 4-35
4-18 Freeing all transport OAM resources 4-37
4-19 Displaying resource configuration information 4-38
4-20 Resolving resource configuration validation errors 4-39
4-21 Addressing classifier resource allocation too small for current configuration error 4-40
4-22 Displaying resource configuration in the configuration file 4-42

Synchronization configuration 5-1

Synchronous Ethernet 5-1
Clock selection algorithm 5-2
Building Integrated Timing Supply 5-3
Network configuration examples 5-3

Copyright© 2012-2013 Ciena® Corporation

List of procedures
5-1 Enabling and disabling synchronization 5-5
5-2 Configuring synchronization 5-6
5-3 Configuring frequency input references 5-8
5-4 Configuring frequency output references 5-12
5-5 Configuring frequency protection groups 5-15
5-6 Creating an output port for SyncE from a BITS input 5-18
5-7 Creating a BITS output from a SyncE input 5-20
5-8 Displaying information for synchronization 5-22

Link Layer Discovery Protocol (LLDP) configuration 6-1

LLDP TLVs 6-3
Feature Benefits 6-6

List of procedures
6-1 Configuring LLDP 6-7
6-2 Configuring TLV transmission 6-9
6-3 Displaying LLDP neighbors 6-11
6-4 Enabling and disabling SNMP notifications 6-12

VLAN management 7-1

VLANs and traffic flow 7-1
Acceptable Frame Types 7-1
Port VLAN ID 7-2
VLAN Ingress Filter 7-2
Egress Untagged VLAN 7-3
Behavior summary 7-4
VLAN/port configuration 7-4
VLAN translation 7-5

List of procedures
7-1 Changing the TPID stamp for a VLAN 7-7
7-2 Configuring a VLAN/port pair to ingress tagged traffic and egress tagged traffic 7-8
7-3 Configuring a VLAN/port pair to ingress untagged traffic and egress untagged traffic 7-10
7-4 Changing tag status 7-11
7-5 Configuring hybrid traffic 7-13
7-6 Emulating a tagged Ethernet port 7-14
7-7 Translating a single NNI VLAN 7-15
7-8 Translating a dual NNI VLAN 7-17

IP management 8-1

IPv6 8-1
IPv6 address format 8-1
IP address usage 8-2
EVC ping 8-5

List of procedures
8-1 Creating interfaces 8-8

Copyright© 2012-2013 Ciena® Corporation

8-2 Deleting an IP or loopback interface 8-10
8-3 Modifying an IP or loopback interface 8-11
8-4 Disabling an IP interface 8-12
8-5 Enabling an IP interface 8-13
8-6 Configuring IPv6 interfaces manually 8-14
8-7 Displaying an IP interface 8-16
8-8 Adding an IP route 8-20
8-9 Removing an IP route 8-21
8-10 Displaying the routing table 8-22
8-11 Displaying FIB entries 8-23
8-12 Enabling Layer 3 switching 8-25
8-13 Disabling Layer 3 switching 8-26
8-14 Displaying the status of Layer 3 switching 8-27
8-15 Clearing all FIB or AIB entries 8-28
8-16 Enabling logging of FIB or AIB events 8-29
8-17 Displaying FIB information 8-30
8-18 Adding a static ARP entry 8-31
8-19 Removing static ARP entries 8-32
8-20 Configuring EVC ping for use case A 8-33
8-21 Configuring EVC ping for use case B 8-36
8-22 Configuring EVC ping for use case C 8-39
8-23 Configuring EVC ping for use case D 8-42
8-24 Configuring EVC ping for use case E 8-46
8-25 Displaying static ARP entries 8-50
8-26 Displaying AIB information 8-51

MEF L2 VPN configuration 9-1

Virtual switch 9-2
Private forwarding groups 9-3
Policy sets 9-4
Port membership 9-5
Virtual circuits 9-6
Provider VLAN 9-7
Reserved VLAN 9-7
Q-in-Q encapsulation 9-8
Q-in-Q configurations 9-8
Frame flooding behavior 9-9
Q-in-Q Ethertype 9-9
EVPL CoS 9-10
Virtual Switch CoS Policies 9-10
VS Member CoS Policy Override 9-10
EVPL Bundling 9-11
VLAN translation for virtual switches 9-12
Ingress push; egress pop (i-push,e-pop) 9-12
Ingress push; egress pop and stamp (i-push,e-pop:stamp) 9-13
Ingress stamp; egress match, pop, and stamp (i-stamp:push,e-match-pop:stamp) 9-13

List of procedures
9-1 Creating an EPL provider bridge 9-16

Copyright© 2012-2013 Ciena® Corporation

9-2 Creating an EVPL provider bridge 9-18
9-3 Configuring the fixed encapsulation priority value 9-20
9-4 Setting the CoS policy when adding VS members 9-21
9-5 Handling ingress untagged frames 9-22
9-6 Setting the L2 transform action on a port 9-24
9-7 Creating an i-push, e-pop Q-in-Q VS configuration 9-25
9-8 Creating a VS configuration with UNI only with bundled CVIDs 9-27
9-9 Creating an i-push, e-pop: stamp configuration 9-29
9-10 Creating an i-stamp:push,e-match-pop:stamp configuration 9-32
9-11 Configuring PFGs 9-35
9-12 Disabling the PFG feature 9-39
9-13 Displaying the configuration of EVPL VS members 9-40
9-14 Displaying virtual switches 9-41
9-15 Displaying PFG information 9-42

Provider Backbone Bridge Traffic Engineering (PBB-TE) implementation 10-1

VLAN Tagging 10-2
Q-in-Q 10-2
MAC Header Encapsulation 10-3
PBB-TE Tunnels 10-4
Connectivity Fault Management 10-7
PBB-TE Dual Homing 10-9
Tunnel Pairing and Synchronization 10-9
Benefits 10-10

List of procedures

10-1 Verifying that a port can participate in PBB-TE 10-12
10-2 Switching from non-native to native PBB-TE support 10-13
10-3 Enabling tunnel synchronization 10-14
10-4 Disabling tunnel synchronization 10-15
10-5 Configuring PBB-TE 10-16
10-6 Releasing reserved BVIDs 10-34
10-7 Displaying PBB-TE information 10-35

Multiprotocol Label Switching (MPLS) Layer 2 VPN configuration 11-1

Terms and definitions 11-2
System requirements 11-3
MPLS services 11-3
VPWS 11-3
VPLS 11-4
H-VPLS 11-5
VPLS membership and MAC learning 11-6
Remote Management for MPLS 11-6
Routing protocols 11-8
OSPF 11-9
IS-IS 11-10
Signaling protocols 11-11

Copyright© 2012-2013 Ciena® Corporation

RSVP-TE 11-11
LDP 11-12
Connectivity Fault Management over MPLS 11-12
Bidirectional Forwarding Detection (BFD) 11-12
Complementary protocols 11-12
LSP ping 11-13
LSP traceroute 11-13
Benefits 11-13
Vendor interoperability 11-13
Configuring L2 VPNs 11-13
Virtual circuits 11-14
Primary virtual circuit attributes 11-14
Comparing raw and tagged PW type for virtual circuits 11-15

List of procedures

11-1 Building an OSPF configuration 11-17
11-2 Configuring an OSPF area 11-19
11-3 Displaying OSPF information 11-21
11-4 Modifying OSPF attributes for an IP interface 11-24
11-5 Clearing OSPF authentication for an OSPF IP interface 11-26
11-6 Displaying OSPF interface information 11-27
11-7 Creating an IS-IS instance with a single area 11-28
11-8 Deleting an IS-IS instance 11-29
11-9 Configuring IS-IS areas 11-30
11-10 Enabling or disabling an IS-IS instance 11-31
11-11 Displaying IS-IS instance protocol information 11-32
11-12 Clearing IS-IS instance statistics 11-33
11-13 Configuring IS-IS 11-34
11-14 Displaying IS-IS interface authentication 11-39
11-15 Associating an IP interface with an IS-IS instance 11-40
11-16 Disassociating an IP interface from an IS-IS instance 11-41
11-17 Enabling or disabling IS-IS for an IP interface 11-42
11-18 Configuring IS-IS padded hello for an IP interface 11-43
11-19 Configuring IS-IS interface level attributes 11-44
11-20 Enabling or disabling RSVP-TE 11-45
11-21 Configuring global RSVP-TE attributes 11-47
11-22 Displaying RSVP-TE configuration and statistics 11-48
11-23 Configuring RSVP-TE attributes for IP interfaces 11-49
11-24 Configuring RSVP-TE IP address authentication 11-50
11-25 Configuring RSVP-TE paths 11-52
11-26 Displaying all or specific paths 11-53
11-27 Running traceroute and ping for RSVP-TE tunnels 11-54
11-28 Configuring tunnels 11-55
Creating and modifying ingress tunnels 11-56
11-29 Displaying tunnel bandwidth profiles 11-59
11-30 Configuring dynamic ingress tunnels 11-60
11-31 Configuring static ingress tunnels 11-61
11-32 Configuring static egress tunnels 11-62
11-33 Configuring static transit tunnels 11-63
11-34 Enabling and disabling tunnels 11-64

Copyright© 2012-2013 Ciena® Corporation

11-35 Deleting ingress tunnels 11-65
11-36 Displaying tunnels 11-66
11-37 Configuring LDP globally 11-67
11-38 Displaying LDP configuration 11-68
11-39 Configuring LDP authentication 11-69
11-40 Configuring virtual circuits 11-70
11-41 Displaying virtual circuits 11-72
11-42 Disabling and enabling virtual circuits 11-73
11-43 Deleting a virtual circuit 11-74
11-44 Configuring label ranges 11-75
11-45 Displaying label ranges 11-77
11-46 Base MPLS configuration (LSR) example 11-78
Example 11-80
11-47 Allocating resources for an MPLS management VS (3916, 3930 and 3931 platforms) 11-81
11-48 Creating an MPLS management VS 11-83
11-49 Displaying remote interface configuration 11-85
11-50 Deleting an MPLS management VS 11-86
11-51 VPWS configuration example 11-87
11-52 VPLS configuration example 11-91
11-53 H-VPLS Configuration 11-94

L2 Control Frame Tunneling configuration 12-1

Overview 12-1
Tunnel method 12-6
Configuration 12-8

List of procedures

12-1 Adding and removing untagged L2 control frame classification 12-9
12-2 Enabling and disabling L2 control frame tunneling 12-10
12-3 Adding and removing control protocols 12-11
12-4 Setting the disposition of control protocols 12-13
12-5 Setting Layer 2 CFT attributes 12-14
12-6 Displaying enabled L2 control frame tunneling instances 12-15
12-7 Displaying the L2 control frame classification for a port 12-17
12-8 Displaying L2 control frame tunneling configuration in the configuration file 12-18
12-9 Configuring L2 Control Frame Tunneling 12-19

Quality of Service configuration 13-1

Class of Services (CoS) policies and mapping 13-2
Traffic profiling 13-4
Determining classification attributes 13-4
Determining Metering Attributes 13-6
Traffic profiling with hierarchical ingress metering 13-7
Advanced classification mode 13-8
Congestion management 13-10
Creating and modifying sRED profiles 13-10
Creating and modifying sWRED profiles on the 3960 13-11
Creating and modifying sWRED profiles on the 3916, 3930, 3931, and

Copyright© 2012-2013 Ciena® Corporation

5150 13-12
Egress scheduling 13-14
Egress shaping 13-15

List of procedures

13-1 Configuring Class of Services (CoS) policies on a port 13-17
13-2 Displaying the CoS policy and mapping on a port 13-19
13-3 Creating a resolved CoS map 13-20
13-4 Modifying a resolved CoS map 13-22
13-5 Setting the resolved CoS map for a port 13-23
13-6 Deleting a custom R-CoS map 13-24
13-7 Displaying resolved CoS maps 13-25
13-8 Creating a frame CoS map 13-27
13-9 Modifying a frame CoS map 13-29
13-10 Deleting a frame CoS map 13-30
13-11 Setting the frame CoS map for a port 13-31
13-12 Displaying frame CoS maps 13-32
13-13 Configuring ingress R-CoS to egress queue mapping 13-33
13-14 Displaying an R-CoS map 13-34
13-15 Applying R-CoS policies and mapping in a VLAN 13-35
13-16 Applying R-CoS policies and mapping in a virtual switch 13-37
13-17 Enabling traffic profiling 13-39
13-18 Setting the traffic profiling provisioning mode 13-40
13-19 Displaying traffic profiling information 13-41
13-20 Setting traffic profiling port attributes 13-43
13-21 Configuring a traffic profiling standard profile 13-47
13-22 Configuring per-port standard traffic profiling 13-53
13-23 Configuring per-port and per-VLAN standard traffic profiling 13-54
13-24 Configuring hierarchical VLAN traffic profiles 13-56
13-25 Configuring hierarchical port traffic profiles 13-61
13-26 Configuring VS classification for standard traffic profiles 13-63
13-27 Configuring VS classification for HIM traffic profiles 13-64
13-28 Displaying standard traffic profiles 13-65
13-29 Clearing statistics for all standard profiles 13-67
13-30 Displaying statistics for all standard profiles 13-68
13-31 Displaying throughput statistics for a traffic profile 13-69
13-32 Setting the per-port hierarchical traffic-profiling mode 13-70
13-33 Setting the child mode 13-71
13-34 Creating an sRED profile 13-72
13-35 Modifying an sRED profile 13-73
13-36 Creating an sWRED profile 13-74
13-37 Modifying an sWRED profile 13-75
13-38 Creating an sWRED profile 13-76
13-39 Modifying an sWRED profile 13-77
13-40 Displaying custom congestion avoidance profiles 13-78
13-41 Updating the congestion avoidance profile for an egress port queue 13-80
13-42 Clearing the congestion avoidance profile to the default for an egress port queue 13-81
13-43 Displaying the congestion avoidance profile for an egress port queue 13-82
13-44 Deleting a custom congestion avoidance profile 13-84

Copyright© 2012-2013 Ciena® Corporation

13-45 Renaming custom congestion avoidance profile 13-85
13-46 Changing the algotrithm of an egress port scheduler 13-86
13-47 Changing the weight of the scheduler for a queue 13-87
13-48 Displaying queue weight and scheduler algorithms 13-88
13-49 Configuring egress port and queue shaping 13-89
13-50 Displaying egress port queue configuration 13-92
13-51 Displaying egress port queue statistics 13-94
13-52 Configuring frame bandwidth calculation 13-95

Multicast services configuration 14-1

IGMP snooping 14-1
- IGMP message shaping 14-1
- Proxy query delay 14-2
- IGMP query engine 14-2
- IGMP message filtering 14-3
- StreamCast 14-5
- IGMP leave modes 14-6
- IGMP parameters 14-7

Unknown multicast filtering 14-8
Configuring multicast 14-9

List of procedures

14-1 Configuring a multicast router topology 14-10
14-2 Configuring a multicast server topology 14-13
14-3 Configuring multicast servers with redundant routers 14-16
14-4 Configuring redudant query engines 14-19

Appendix A: Error codes A-1

Traffic profiling error codes A-1

Copyright© 2012-2013 Ciena® Corporation

Copyright© 2012-2013 Ciena® Corporation

About this document

This manual describes how to configure system software on 39XX/51XX Service Delivery and Aggregation switches. This system software is based on a common Service Aware Operating System (SAOS) code base designed to deliver consistent benefits across all Ethernet delivery, aggregation, and distribution configurations.

> ***Note:*** This system software cannot be installed on any other Service Delivery Switches, Service Concentration Switches or Service Aggregation Switches.

This manual provides information and examples for use in configuring system software on any platform on which it is installed. It includes an explanation of the key features supported by the devices and provides example configurations for these features. Although these examples are useful in configuration, they are not meant to be used as a configuration template.

Conventions used in this document

Hyperlinks are indicated by blue text in this document.

In procedures, the following text conventions are used:

- `courier` text, for system responses
- *italic* text, for expected results
- **bold** text, for user input

Copyright© 2012-2013 Ciena® Corporation

Command syntax

A variety of symbols are used to indicate CLI command syntax. These symbols describe how to enter a command. They are not entered as part of the command itself. The following table summarizes command syntax symbols.

Symbol	Description
< >	Encloses a variable or literal value that must be specified. Some examples include: server <IpAddress> priority <NUMBER: 1-7> dns <on|off> description <String[31]> For server <IpAddress>, the attribute could be entered as server 10.10.11.100 or server www.ciena.com. With priority <NUMBER: 1-7> the text within <> indicates that 1 - 7 are valid values. In the example of dns <on|off>, either the literal value of on or off is valid, such as dns on. For description <String[31]>, any string of up to 31 characters is entered.
{ }	Encloses a required value or list of required arguments. One or more values or arguments can be specified. For example, in the syntax: cfm mip create {vlan <VlanId>} {port <PortNameList>} [level <NUMBER: 0-7>] The vlan and port arguments are required. The level argument is optional.

|	Separates mutually exclusive items in a list, only one of which can be entered. For example, in the syntax:
	`dhcp client options set subnet <on|off>`
	Either on or off must be specified, for example:
	`dhcp client options set subnet on`

[ ]	Encloses an optional value or a list of optional arguments. One or more values or arguments can be specified. For example, in the syntax:
	`arp show [interface <Interface>]`
	You can enter a value for interface <Interface> or not. For example:
	`arp show`

Copyright© 2012-2013 Ciena® Corporation

Symbol	Description
{ [ ], [ ], [ ] }	Specifies a list of optional items where at least one must be specified.
…	Indicates the example has been abbreviated and that the actual display contains more information.
*	Indicates zero or more occurrences of what is preceding.

New in this release

For information about what is new in SAOS 6.10, see *39XX/51XX R6.10 Release Notes* (009-3203-013).

Documents in the 39XX/51XX documentation suite

For descriptions of documents in the 39XX/51XX Service Delivery and Aggregation Switches documentation suite, see *39XX/51XX Service Delivery and Aggregation Switches Product Fundamentals* (009-3203-006).

Copyright© 2012-2013 Ciena® Corporation

Copyright© 2012-2013 Ciena® Corporation

Configuration fundamentals

This chapter provides an overview of the following components for building networks and accounting:

- "Ports"
- "Hardware resources"
- "Synchronization configuration"
- "Link Layer Discovery Protocol"
- "Virtual Local Area Networks"
- "IP management"
- "Metro Ethernet Forum Layer 2 Virtual Private Networks"
- "Provider Backbone Bridge Traffic Engineering"
- "Multiprotocol Label Switching"
- "Layer 2 control frame tunneling"
- "Quality of Service"
- "Multicast services"

Ports

Physical ports provide connectivity to other devices. Logical ports are created when multiple physical ports are joined in a Link Aggregation Group (LAG).

Physical ports provide connectivity to other devices, which is essential for any switching device. To aggregate bandwidth and provide link redundancy between two devices, physical ports are added to a Link Aggregation Group (LAG). The port management commands provide the ability to configure ports and troubleshoot connectivity.

Hardware resources

The system assigns hardware resources (classifier, meter, and counter resource types) for various software features. Depending upon the feature, you can reassign these resources to provide additional resources for other features.

Copyright© 2012-2013 Ciena® Corporation

This customized resource management is supported for the following features:

- Broadcast containment
- CFM
- DHCP relay
- Traffic profiling
- Virtual circuit statistics
- Virtual switch Layer 2 enhanced transforms
- Transport OAM

Synchronization configuration

System timing recovers and distributes frequency, phase and time-of-day (ToD) information in order to maintain synchronization between network elements.

System timing comprises:

- Synchronous Ethernet (SyncE)
- Building Integrated Timing Supply (BITS)

SyncE assures that frequency is distributed at the physical layer where it is not subject to load impairments such as packet congestion and loss.

BITS is a timing signal that is used to distribute frequency synchronization in a telecom environment, typically in a central office. It is usually carried over T1/E1 lines. Timing is encoded within the transmitted data signal to synchronize the whole network.

Link Layer Discovery Protocol

Link Layer Discovery Protocol (LLDP) allows network equipment, for example, stations, switches, bridges, routers, to advertise their parameters for network topology discovery and management. Traditional network management protocols, such as SNMP, running at key locations, use layer 3 protocols to identify the devices connected to the network. The Link Layer Discovery Protocol is a layer 2 protocol, allowing precise discovery of the physical-link topology of the network. Devices act as LLDP agents, which drastically increases the network discovery performance of SNMP applications, as well as any system capable of accessing standard LLDP MIBs.

Virtual Local Area Networks

Virtual Local Area Networks (VLANs) are used to group resources that have a common set of requirements, regardless of where they are located physically. They also allow ports on a device to be grouped in order to limit the

Copyright© 2012-2013 Ciena® Corporation

distribution of unicast, multicast, and broadcast traffic. For example, flooded traffic originating from a particular VLAN can be limited to only other ports belonging to that VLAN. VLANs allow traffic on the same physical connection to be divided into separate services. These services can then be further divided into groups within each service.

Note: VLAN configuration is supported for compatibility with Enterprise VLAN applications. For Carrier Ethernet applications, use virtual switch configuration as described in "MEF L2 VPN configuration" on page 9-1.

IP management

IP management allows you to manage the switch using IPv4 and/or IPv6. IP management applications are:

- file transfer (xftp)
- software upgrade
- SNMP
- syslog
- RADIUS
- TACACS+
- NTP
- SSH
- Telnet

Metro Ethernet Forum Layer 2 Virtual Private Networks

Layer 2 (L2) Virtual Private Networks (VPNs) support the three Ethernet Service Types defined by the Metro Ethernet Forum (MEF) as:

- E-Line: point-to-point service.
- E-LAN: multiple point service.
- E-Tree: point-to-multipoint service.

These services are defined by the configuration of port and VLAN based Ethernet Private Line/LAN (EPL) or Ethernet Virtual Private Line (EVPL). The customer perspective is that their connection to the service provider is a direct connection to a private line or LAN between their sites.

L2 VPNs transport Ethernet/802.3 and VLAN tagged traffic between multiple sites that belong to the same L2 broadcast domain. These services are deployed over Ethernet using an extension of 802.1Q, called Q-in-Q.

Copyright© 2012-2013 Ciena® Corporation

Provider Backbone Bridge Traffic Engineering

Provider Backbone Bridge Traffic Engineering (PBB-TE) allows providers to create point-to-point, primary and backup Ethernet tunnels and specify the path that traffic will take across their Ethernet metro networks. These paths reserve appropriate bandwidth and support the provisioned QoS metrics. However, PBB-TE is unique in that it actually disables some Ethernet features in order to accomplish its goal of delivering traffic.

Multiprotocol Label Switching

Multiprotocol Label Switching (MPLS) is a protocol that allows different networks to converge and appear to be a single autonomous network. MPLS can be used to build seamless and consistent Layer 2 (L2) Virtual Private Networks (VPNs) between two or more remote sites traversing many different physical networks. The customer perspective is that their connection to the service provider is a direct connection using a private Local Area Network (LAN) between geographically separated sites.

Layer 2 control frame tunneling

A device identifies the associated protocol of L2 control frame based upon the Media Access Control Destination Address (MAC DA) and other information within the frame. Depending upon the state of processing, L2 control frames can be in three forms:

- untagged: standards-based definition of the protocol's control frame. In this form, the frame does not have a VLAN tag, but does have the protocol specific MAC DA. Depending upon the protocol, it can have a specific Ethertype value and other information. Typically, control frames in this format are received from the subscriber facing interface.
- transparent (tagged): one or more 802.1 VLAN tags have been added to transform the frame, and the protocol specific MAC DA is left intact. This form occurs when a control frame is received from a subscriber facing interface, encapsulated as a data frame, and then forwarded from a network facing interface for transport through the provider network.
- L2 Protocol Tunneling (L2PT): standard MAC DA has been transformed to the L2PT special MAC DA.

Copyright© 2012-2013 Ciena® Corporation

Quality of Service

Quality of Service (QoS) refers to the management of bandwidth to ensure that network traffic is allocated the desired amount of network resources. Table 1-1 lists mechanisms for managing bandwidth.

Table 1-1
Mechanisms for managing bandwidth

Mechanism	Description
Class of Services (CoS) policies and mapping	CoS policies and mapping comprise the following: • Resolved CoS: applies a Resolved CoS Policy and Resolved CoS Map on ingress. Optionally, remarks the frame's Layer 2 priority and color based upon the mapping. Assigns traffic to egress CoS queues based upon Resolved CoS values. Ingress coloring influences both traffic profiling and congestion avoidance processing. Traffic profiling meters are now "color-aware" as of release 6.8.0. These meters respond to ingress R-Color, but do not respond to ingress R-CoS mappings. In release 6.9.1 and later, you can configure whether a meter is color blind or color aware per traffic profile. • Egress Frame CoS Policy: enables an egress frame CoS policy. ***Note:*** This setting is not available on 3902, 3911, 3920, 3940, and 5140 devices. • Frame CoS Mapping - applies a R-CoS and R-Color to frame Priority Code Point (PCP)/Layer 2 (L2) CoS 802.1D priority and Discard Eligibility Indicator (DEI)/Canonical Format Indicator (CFI) mapping that remarks egressing L2 frames.
Traffic profiling and Traffic profiling with hierarchical ingress metering	Provides ingress traffic classification and metering. ***Note:*** To configure traffic profiling, you need to install the Advanced Ethernet license key. To obtain the Advanced Ethernet license key, contact Ciena Sales.
Congestion management	Method for managing CoS queue traffic when congestion occurs on egress. ***Note:*** To enable configurable congestion avoidance, you need to install the Advanced Ethernet license key. To obtain the Advanced Ethernet license key, contact Ciena Sales.
Egress scheduling	Determines the order in which the physical queues are processed.
Egress shaping	Controls bandwidth for taking frames out of queues at egress
Configurable frame bandwidth calculation	Configure whether to use the inter-frame-gap (IFG) in the calculations for broadcast containment, ingress metering and egress shaping.

Multicast services

Traditional Internet Group Management Protocol (IGMP) was designed for environments that have a low volume of multicast packets and no real-time traffic requirements for processing IGMP messages. Ciena's IGMP implementation was designed for networks where multicast services are critical, such as networks delivering IP broadcast video. Devices employ enhanced IGMP snooping, Channel Streams, and various filters to limit multicast streams and assure their timely delivery.

The following network elements support 256 multicast groups:

- 3902
- 3911
- 3920

The following network elements support 1024 multicast groups:

- 3916
- 3930
- 3931
- 3940
- 3960
- 5140
- 5150

Copyright© 2012-2013 Ciena® Corporation

Configuration management

This chapter provides explanations for implementing the basic set up of a device and describes the various options available to configure it, including:

- "Accessing the CLI" on page 2-1
- "Configuration files" on page 2-1

This chapter provides the procedures for configuration management.

Accessing the CLI

To access CLI commands, connect to the device by establishing a Telnet session or through direct connection to the serial console port located on the front of the control module.

For related procedures, see *39xx/51xx R6.10 Administration and Security* (009-3203-007).

Configuration files

A device can store multiple device configuration files. However, only one configuration file can be active at a time. By default, configuration information is saved to a file called `startup-config`. The `startup-config` file is also the default load file. The parameters defined in the `startup-config` file are applied when the device reboots (unless an alternate file is specified). The current running configurations on a device are not saved to a configuration file unless specifically saved. This includes configuration changes made using the CLI or SNMP. If a device is rebooted without saving the configuration, all changes are lost.

This chapter provides the following procedures:

- "Saving configuration changes" on page 2-3
- "Displaying configuration files" on page 2-4
- "Augmenting the current configuration" on page 2-5
- "Restoring default configurations" on page 2-7
- "Setting the default configuration files" on page 2-8

Copyright© 2012-2013 Ciena® Corporation

- "Displaying the default configuration files" on page 2-9
- "Resetting default configuration files to factory default files" on page 2-10

Copyright© 2012-2013 Ciena® Corporation

Procedure 2-1
Saving configuration changes

In order to permanently save configuration changes, you must save the running configuration to a configuration file. To save the running configuration, use the `configuration save` command. By default, the command saves the current configuration to the default configuration file, `startup-config`.

By saving alternate versions of command files, you can store multiple configuration files for running different configurations of the system. For example, you can save configuration to an alternate file as a backup to restore to a previous configuration or you can store a configuration file for configuring another device of the same family.

> ***Note:*** If the running configuration has not been saved, an asterisk is displayed in the CLI prompt, such as: `*>`. When you save the configuration, the asterisk is removed.

Step	Action

To save the running configuration

1 Save the running configuration:

```
configuration save
```

To save to an alternate filename

2 Save the configuration to an alternate filename:

```
configuration save filename <filename>
```

where

filename <filename> is the configuration file name.

—end—

Example

The following example saves the default configuration file.

```
configuration save
```

The following example saves to an alternate filename.

```
configuration save filename myConfig
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 2-2
Displaying configuration files

Display configuration files.

Step	Action

1 Display the configuration:

```
configuration show [line-numbered] [differences-from-saved]
```

where

line-numbered	displays with line numbers.
differences-from-saved	displays changes since last saved.

2 Search the configuration file:

```
configuration search file <String>] [lines <NUMBER: 0-40>] {string <String>}
```

where

file <String>]	is the filename (no path).
lines <NUMBER: 0-40>	is the search window size.
string <String>	is the search string.

—end—

Example

The following example displays the differences between the running configuration and the saved configuration file.

```
> configuration show differences-from-saved
diff /flash0/config/Mcast_Aggs_DG /ram/65700.out
4,6c4,6
< ! Created:       Mon May  5 14:04:02 2008
---
> ! Created:       Mon May  5 14:09:04 2008
```

The following example displays sections of a configuration file containing a specific string.

```
> configuration search string dhcp
dhcp client disable
dhcp client set interface local
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 2-3
Augmenting the current configuration

You can add configuration commands from a file to the current configuration with the `configuration augment` command. By default, the commands are added from a specified file stored on the file system at /flash0/config.

Note: If you want to apply the added commands to the default startup command file, you need to save the configuration with the `configuration save` command.

Step	Action

1 Augment the current configuration with a file:

```
configuration augment {filename <String>} [server <IP
address or host name>] {default-tftp-server} {default-
ftp-server} {default-server} {default-sftp-server}
[sftp-server <IP address or host name>] [tftp-server <IP
address or host name>] [ftp-server <IP address or host
name>] [login-id <String[32]>] [password <Password
String>] [secret <String[256]>]
```

where

filename <String>	is the configuration file name.
server <IP address or host name>	is the FTP server.
default-tftp-server	indicates that the default TFTP server is used.
default-ftp-server	indicates that the default FTP server is used.
default-server	indicates that the default xFTP server is used.
default-sftp-server	indicates that the default SFTP server is used.
sftp-server <IP address or host name>	is the SFTP .
tftp-server <IP address or host name>	is the TFTP server.
ftp-server <IP address or host name>	is the FTP server.

Copyright© 2012-2013 Ciena® Corporation

where

login-id <String[32]>	is the FTP/SFTP username.
password <Password String>	enter the password in clear text.
secret <String[256]>	sets the password using a pre-encrypted string.

—end—

Example

The following example adds configuration commands to the configuration.

```
configuration augment filename /users/testuser/
myConfigAdd.txt default-server
WORKING: TFTP file transfer in progress
/users/testuser/myConfigA 100%
|**************************************************|
70   0:00:00 ETA
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 2-4
Restoring default configurations

You can

- reset configuration to user defaults
- reset the device to factory defaults

After running any of the commands to activate an alternative configuration, you have 10 seconds to cancel. To cancel, press Ctrl+C.



CAUTION
Loss of Configuration Information
When you reset a device to its factory default settings, all configuration and file system changes are lost, including saved configuration files and log files. Software License Keys that were previously installed are not removed from the device.

Step	Action

To reset configuration to user defaults

1 Reset configuration to user defaults:

```
configuration reset-to-user-config [filename <filename>]
```

where

filename <filename> is the configuration file name.

To restore the device to factory defaults

2 Restore the device to factory defaults:

```
configuration reset-to-factory-defaults
```

—end—

Example

The following example resets the configuration to user defaults.

```
configuration reset-to-user-config filename myConfig
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 2-5
Setting the default configuration files

You can set alternate configuration files as the default save file so that when a configuration save is performed, the changes will be saved to the alternate configuration file. When the device is rebooted, it will load the default-load file, which may or may not be the same file as the default-save file. However, when the device is rebooted it will load the startup-config file.

You can

- set the default file for saving configuration
- set the default file for loading configuration

Step	Action

To set the default file for saving configuration

1 Set the default file for saving configuration:

```
configuration set default-save-filename <FileName>
```

To set the default file for loading configuration

2 Set the default file for loading configuration:

```
configuration set default-load-filename <FileName>
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 2-6
Displaying the default configuration files

The default configuration files are

- save
- load

Step	Action
1	Display the default save, load, and backup load files: `configuration list`

Example

The following example shows sample output for the configuration list command.

```
> configuration list

+------------------------------------------------------------------------------+
|      Configuration Files                                                     |
+------------------------------------------------------------------------------+
| startup-config                                                               |
| test                                                                         |
+------------------------------------------------------------------------------+
| Default Save File: test                                                      |
| Default Load File: test                                                      |
+------------------------------------------------------------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 2-7
Resetting default configuration files to factory default files

You can

- reset the default file for saving configuration
- reset the default file for loading configuration

Step	Action

To reset the default file for saving configuration

1 Reset the default file for saving configuration:

```
configuration unset default-save-filename
```

To reset the default file for loading configuration

2 Reset the default file for loading configuration:

```
configuration unset default-load-filename
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Port management

This chapter explains how to configure physical and logical port attributes.

Physical ports provide connectivity to other devices, which is essential for any switching device. To aggregate bandwidth and provide link redundancy between two devices, physical ports are added to a Link Aggregation Group (LAG). The port management commands provide the ability to configure ports and troubleshoot connectivity.

Port management addresses:

- "Port attributes"
- "Port statistics"
- "Transceivers"

This chapter provides the procedures for port management.

Port attributes

Table 3-1 describes administrative and operational attributes for ports.

Table 3-1
Administrative and operational attributes for ports

Attribute	Description
General	
Port Name	A 32 character string representing the name of the physical port or LAG. For physical ports, the name represents the port's physical location identifying the chassis module and port in the format: <ModuleNumber>.<PortNumber> Whenever a platform has a single module, the component number will be left out of the name. For example, the 3960 platform is single module, so each physical port is named only with the port number (1 through 12). The 5150 platform supports multiple modules, so each port on the second and third module is named with the module and port numbers (2.1, 2.2, 3.1, 3.2).
Type	Gigabit (GIG), 10 Gigabit (10GIG), or LAG.

Copyright© 2012-2013 Ciena® Corporation

Table 3-1
Administrative and operational attributes for ports

Attribute	Description
Service Port Type	Indicates whether the port is a Subscriber, User to Network Interface (UNI), or Network to Network Interface (NNI) port.
Spanning Tree State	Indicates the state of Spanning Tree Protocol (STP), including: Disabled, Forwarding, Learning, or Discarding. For additional information regarding STP, refer to “Configuring RSTP” on page 19-1 and “Configuring MSTP” on page 20-1.
MAC Address	Media Access Control (MAC) address. By default, the MAC address is uniquely assigned during manufacturing.
Link Flap Detection	configuration for monitoring link state transitions.
Link State	Indicates whether the port is enabled or disabled. By default the Admin link state is disabled.
State Group Link State	Indicates whether the port state mirror group link is enabled or disabled. Blank indicates the port is not a member of a port state mirror group.
Acceptable frame type (acceptable-frame-type)	Designates the treatment of received frames to allow all, tagged-only, or untagged-only.
Flow Control (flow-ctrl)	Prevents one port from sending data faster than the receiving port can handle it. When the receiving port has all the data it can handle, it sends a “pause” frame to the sender. The sender stops sending data until the pause frame expires. Received (asym-rx), transmitted only (asym-tx), or Off modes are supported; the default mode is off.
Auto-negotiation (auto-neg)	Determines whether ports negotiate with their link partner to operate with parameters common to both links. This method of auto negotiation follows the IEEE 802.3z standard and provides a way to automatically connect multiple types of devices. By default, auto negotiation is enabled.
Flow Control Advertised (advertised-flow-control)	Determines whether flow control setting is advertised. Default is off.
Duplex (duplex)	Half or full. When the port is set to full duplex, it can transmit and receive data simultaneously. With half duplex the port can transmit or receive data, but not both simultaneously. Default duplex is set to full.
Description (description)	Configurable 128 character description of the port. By default, the description is blank.

Copyright© 2012-2013 Ciena® Corporation

Table 3-1
Administrative and operational attributes for ports

Attribute	Description
Speed	Physical port speed, such as 1 Gbps or 10 Gbps. Not applicable for LAG ports. The default value is auto, which matches the speed to the transceiver speed. Any configured auto negotiation settings are ignored for tranceivers that do not support autoneg, that is, 100M- and 10G-based transceivers. ***Note:*** If you set a value for the speed attribute, the port stays in that speed and a transceiver mismatch error is displayed if there is a mismatch. This functionality is only supported on Ciena-supported transceivers.
Maximum frame size (max-frame-size)	Maximum frame size in bytes allowed to ingress the port. The default value is 1526. Jumbo frames are supported with configurable range from 1522-9216. Maximum frame size is also referred to as Maximum Transmission Unit (MTU) size.
Aggregation Membership	Displays the link aggregation group of which the port is a member.
Port traffic mirroring	
Mirror port	Turns port mirroring on or off. Default is off.
Ingress-mirror	Sets the mirror port for ingress traffic.
Egress-mirror	Sets the mirror port for egress traffic.
Optic transceiver	
Mode	Shows the port connector mode, Copper RJ45, Small Form Factor Pluggable (SFP), or SFP+.
XCVR Capabilities Mismatch	Indicates whether or not the capabilities of the port and the installed transceiver match.
Phy loopback	
Loopback	Indicates whether internal physical loopback is enabled. By default, internal loopback is off. Internal loopback is supported where data that is received on a port in internal loopback mode will egress the same port, and the loopback occurs in the PHY. Internal loopback can be enabled on any physical port independently, regardless of VLAN membership. Setting the internal loopback attribute automatically sets the port's learn limit to 0, with an action of 'forward'. When the internal loopback setting is 'off', the configured learn limit is then re-applied. If more than two ports within the same VLAN are configured to participate in an internal loopback test, there is a danger of creating a broadcast storm.
Class of Service (CoS)	

Copyright© 2012-2013 Ciena® Corporation

Table 3-1
Administrative and operational attributes for ports

Attribute	Description
Fixed Resolved CoS (fixed-rcos)	Sets the fixed resolved CoS value.
Fixed Resolved Color (fixed-rcolor)	Sets the fixed resolved color, which is green or yellow.
Resolved CoS (R-CoS) to Frame CoS (F-CoS) map (frame-cos-map)	Sets the R-CoS to F-CoS map.
Resolved CoS Policy (resolved-cos-policy)	Sets Resolved CoS Policy.
Ingress to Egress Qmap (ingress-to-egress-qmap)	Sets the R-CoS queue map to use in mapping internal CoS (R-CoS) at the ingress port to a CoS queue at an egress port.
FCOS to RCOS Map (resolved-cos-map)	Sets the frame CoS to R-CoS map.
Remark L2 (resolved-cos-remark-l2)	Enables or disables frame layer 2 remarking.
VLAN specific	
Egress Untag VLAN (egress-untag-vlan)	Sets the VLAN for egressing untagged data frames.
PVID (PVID)	Port VLAN ID. Default is 1.
Ingress VLAN Filter (vlan-ingress-filter)	Filters frames that are not members of a configured VLAN. Default is enabled.
VLAN Membership	Displays the VLANs of which the port is a member.
Virtual switch and virtual circuit specific	
Untagged Ctrl VS (untagged-ctrl-vs)	Untagged control frame virtual switch.
Untagged Data VS (untagged-data-vs)	Untagged data frame virtual switch.
Untagged Ingress Data Vid (untagged-data-vid)	Pushes and pops the specified VID as a Customer VID for frames forwarded to a virtual switch. Applicable only to ports associated with virtual switches. ***Note:*** On the 3902, 3911, 3920, 3940 and 5140 platforms, the untagged-data-vid attribute pushes the specified VID as a Customer VID on ingress but does not pop the VID on egress.
Virtual switch ingress filter (vs-ingress-filter)	Filters frames that are not a member of a configured virtual switch.

Copyright© 2012-2013 Ciena® Corporation

Table 3-1
Administrative and operational attributes for ports

Attribute	Description
Eth VC EtherType	Displays the EtherType for Ethernet virtual circuits.
Eth VC EtherType Policy	Displays the configured EtherType policy for the Ethernet virtual circuits.
VS L2 Transform (vs-l2-transform)	Enables VLAN translation for Q-in-Q with virtual switch L2 transform actions.

This chapter provides the following procedures for configuring ports:

- "Setting port attributes" on page 3-12
- "Resetting port attributes to default" on page 3-13
- "Disabling a port" on page 3-14
- "Enabling a port" on page 3-15
- "Displaying port attributes" on page 3-16
- "Displaying blade information" on page 3-28
- "Displaying port capabilities" on page 3-32
- "Displaying port Ethernet configuration" on page 3-34
- "Displaying port status" on page 3-35

Port statistics

Table 3-2 describes port statistics.

Table 3-2
Port statistics

Port statistic	Description
RxBytes	Number of bytes received including those in bad packets.
RxPkts	Number of packets received including all unicast, multicast, broadcast, MAC control and bad packets.
RxCrcErrorPkts	Number of packets received which contained an FCS error and were between 64 and 1518 bytes in length.
RxMcastPkts	Number of good multicast packets received that were between 64 and 15184 bytes in length. Excludes MAC control frames.
RxBcastPkts	Number of good broadcast packets received that were between 64 and 15184 bytes in length. Excludes MAC control frames.
RxUcastPkts	Number of good unicast packets received that were between 64 and Max Frame Size bytes in length. Excludes MAC control frames.

Copyright© 2012-2013 Ciena® Corporation

Table 3-2
Port statistics

Port statistic	Description
UndersizePkts	Number of packets received that were less than 64 bytes long and contained a valid FCS and were otherwise well-formed.
OversizePkts	Number of packets received that were longer than 15184 bytes to Max Frame Size and contained a valid FCS and were otherwise well formed . Includes unicast, multicast and broadcast packets.
FragmentsPkts	Number of packets received that were between 10 and 63 bytes in length and had either an FCS error or an alignment error.
JabbersPkts	Number of packets received that were longer than 15184 bytes to Max Frame Size and had an FCS error or an alignment error.
RxPausePkts	Number of received valid pause packets that were between 64 and 1518 bytes in length.
RxDropPkts	The total number of valid packets received which were discarded due to lack of resources, that is, rx buffer hits the discard limit, buffer pool full or back pressure discard. RFC 2819 specifies that this number is not necessarily the number of packets dropped; it is just the number of times this condition has been detected.
RxDiscardPkts	The total number of valid packets received which were discarded due to lack of resources, that is, rx buffer hits the discard limit, buffer pool full or back pressure discard. RFC 2819 specifies that this number is not necessarily the number of packets dropped; it is just the number of times this condition has been detected.
RxLOutRangePkts	Number of packets received which exceeded Max Frame Size in length and contained a valid or invalid FCS.
64OctsPkts	Number of packets received that were 64 bytes in length.
65To127OctsPkts	Number of packets received that were between 65 and 127 bytes in length.
128To255OctsPkts	Number of packets received that were between 128 and 255 bytes in length.
256To511OctsPkts	Number of packets received that were between 256 and 511 bytes in length.
512To1023OctsPkts	Number of packets received that were between 512 and 1023 bytes in length.
1024To1518OctsPkts	Number of packets received that were between 1024 and 1518 bytes in length.
1519To2047OctsPkts	Number of packets received that were between 1519 and 2047 bytes in length.
2048To4095OctsPkts	Number of packets received that were between 2048 and 4095 bytes in length.

Copyright© 2012-2013 Ciena® Corporation

Table 3-2
Port statistics

Port statistic	Description
4096To9216OctsPkts	Number of packets received that were between 4096 and 9216 bytes in length.
TxBytes	Number of bytes transmitted including those in bad packets.
TxPkts	Number of packets transmitted including all unicast, multicast, broadcast, MAC control and bad packets.
TxExDeferPkts	Number of transmitted packets which experienced multiple deferrals (ie. 2 or more deferrals according to Broadcom documentation).
TxDeferPkts10	Number of transmitted packets which were deferred on the first transmission but did not experience any subsequent collisions during transmission.
TxGiantPkts	Number of packets transmitted that were longer than 15184 bytes2 and were otherwise well formed (valid FCS)
TxUnderRunPkts	Number of transmitted underrun packets.
TxCrcErrorPkts	Number of transmitted packets which contained an FCS error.
TxLCheckErrorPkts	Number of transmitted length check packets
TxLOutRangePkts	Number of transmitted length out of range packets
TxLateCollPkts	Number of transmitted packets which experienced a late collision more than 512 bit times during a transmission attempt
TxExCollPkts	Number of transmitted packets which experienced 16 collisions during transmission and was aborted.
TxSingleCollPkts	Number of transmitted packets which experienced a single collision.
TxCollPkts	Number of transmitted packets which experienced 2-15 collisions (including any late collisions) during transmission.
TxPausePkts	Number of valid pause control packets transmitted that were between 64 and 15184 bytes in length.
TxUcastPkts	Number of good unicast packets transmitted that were between 64 and 15184 bytes in length.
TxMcastPkts	Number of good multicast packets transmitted that were between 64 and 15184 bytes in length.
TxBcastPkts	Number of good broadcast packets transmitted that were between 64 and 15184 bytes in length.
Tx64OcPkts	Number of packets transmitted that were 64 bytes in length.

Copyright© 2012-2013 Ciena® Corporation

Table 3-2
Port statistics

Port statistic	Description
Tx65To127OcPkts	Number of packets transmitted that were between 65 and 127 bytes in length.
Tx128To255OcPkts	Number of packets transmitted that were between 128 and 255 bytes in length.
Tx256To511OcPkts	Number of packets .transmitted that were between 256 and 511 bytes in length.
Tx512To1023OcPkts	Number of packets transmitted that were between 512 and 1023 bytes in length.
Tx1024To1518OcPkts	Number of packets transmitted that were between 1024 and 1518 bytes in length.
Tx1519To2047OcPkts	Number of packets transmitted that were between 1519 and 2047 bytes in length.
Tx2048To4095OcPkts	Number of packets transmitted that were between 2048 and 4095 bytes in length.
Tx4096To9216OcPkts	Number of packets transmitted that were between 4096 and 9216 bytes in length.

This chapter provides the following procedures for statistics:

- “Displaying port statistics” on page 3-17
- “Monitoring port statistics” on page 3-22
- “Clearing current statistics” on page 3-27

Transceivers

This section describes

- “Identification”
- “Diagnostics”

Identification

Ciena devices support transceivers that contain a standard serial erasable programmable read-only memory (EPROM) that provides information on the type of SFP used. The following information is read from the EPROM:

- Identifier Type (GBIC, SFP...)
- Extended Identifier Type
- Connector Type (SC, LC, MU, SG...)
- Vendor Name

Copyright© 2012-2013 Ciena® Corporation

- Vendor Organizational Unique Identifier (OUI)
- Vendor Part Number
- Vendor Serial Number
- Vendor Revision Number
- Encoding Algorithm (NRZ, Manchester...)
- Manufacturing Date Code
- Transceiver Code
- Transceiver SFF-8472 Compliance Version

Diagnostics

Ciena devices support advanced transceivers that have an additional diagnostic serial EPROM. The system software determines if the transceiver has the diagnostic EPROM and will provide the following information to the user:

- Wavelength/Frequency
- Temperature
- Rx Power
- Tx Power
- Tx Disable State
- Tx Fault State
- Rx Rate Select State

The information is stored in a table on a per port basis. The standard EPROM information is updated during initialization or when a new transceiver has been inserted. The diagnostic information is updated at a rate of 1 port per 5 seconds. However, this process has a low priority, and in times of a heavy CPU load, the information may be refreshed slowly or not at all.

Transceivers that support diagnostics can trigger events and SNMP traps.

Diagnostics on SFPs are the following:

- BiasHigh
- BiasLow
- RxPowerHigh
- RxPowerLow
- TempHigh
- TempLow
- TxPowerHigh

Copyright© 2012-2013 Ciena® Corporation

- TxPowerLow
- VccHigh
- VccLow

Diagnostics on XFPs are the following:

- BiasHigh
- BiasLow
- RxPowerHigh
- RxPowerLow
- TempHigh
- TempLow
- TxPowerHigh
- TxPowerLow

Each of these events and traps include a warning and alarm version of each, for example, BiasHighAlarm and BiasHighWarning). Thresholds are set by the SFP vendors, and are not programmable. Both event classes (alarm, warning) are logged under the xcvr-mgr. The warnings are logged using the debug category and warning severity. The alarms are logged using the debug category and minor severity.

These alarms and warnings are based on flags that are set or cleared inside the SFP. These flags are polled at a low priority and slow rate, so flags may be set and then cleared without generating a trap or event. For example, if the TempHigh alarm threshold is exceeded for a few seconds, and then cleared before the flag is polled, it will not trigger a TempHigh alarm. You can forcibly clear alarms and warnings by removing and then reinserting the transceiver or disabling and then enabling the port.

Copyright© 2012-2013 Ciena® Corporation

Table 3-3 lists transceiver states and provides a description of each state.

Table 3-3
Transceiver states

State	Description
INV!	Invalid. The transceiver port state could not be determined due to a system error.
UCTF	Uncertified. The transceiver is not in the officially supported set of transceivers on that device, for that software version. The transceiver may or may not function properly.
WARN	There are one or more warnings associated with the transceiver, such as the port configuration settings (e.g. speed, autonegotiation, etc) do not match the capabilities of the transceiver. For example, the port is configured for Gig speed, but the transceiver is 100m. Information about which settings are incompatible are available in the output of the port show port command.
FLT!	Fault. The transceiver has been faulted for some reason; typically this will be due to EEPROM checksum and/or read failures.
Ena	Enabled
Dis	Disabled

This chapter provides the following procedures for configuring transceivers:

- "Displaying a list of supported optics" on page 3-36
- "Displaying transceiver information" on page 3-37
- "Determining transceiver speed" on page 3-40
- "Tuning XFP transceivers" on page 3-42
- "Setting the port connector mode" on page 3-44

Copyright© 2012-2013 Ciena® Corporation

Procedure 3-1
Setting port attributes

Set port attributes. For information about port attributes, see “Administrative and operational attributes for ports”.

Step	Action
1	Set port attributes: `port set port <PortNameList> {[acceptable-frame-type <all|tagged-only>], [advertised-flow-control <asym-tx|off|sym-asym-rx>],[auto-neg <on|off>], [duplex <half|full>], [description <String[31]>], [egress-frame-cos-policy <ignore|rcos-to-l2-outer-pcp-map>] [egress-untag-vlan <Vlan>], [egress-mirror <PortName>], [fixed-rcos <NUMBER: 0-7], [fixed-rcolor <green|yellow>], [flow-ctrl <asym-rx|asym-tx|off|sym>], [ingress-mirror <PortName>], [link-flap-detect <on|off>], [link-flap-count <#>], [link-flap-detect-time <#>], [link-flap-hold-time <#>], [ingress-to-egress-qmap <RcosQueueMap>] [max-frame-size <NUMBER: 1522-9216>], [mirror-port <on|off>], [mode <default|rj45|sfp>], [loopback <on|off>],[pvid <Vlan>], [resolved-cos-policy <dot1d-tag1-cos|fixed-cos|l3-dscp-cos>], [resolved-cos-map <FcosRcosMap>], resolved-cos-remark-l2 <true|false>], [speed <ten|hundred|gigabit|auto>], [untagged-ctrl-vs <VirtualSwitchName>], [untagged-data-vs <VirtualSwitchName>], [untagged-data-vid <VlanId>], [vlan-ingress-filter <on>], [vs-ingress-filter <on>], [vs-l2-transform <i-push,e-pop|i-push,e-pop:stamp|i-stamp:push,e-match-pop:stamp>]}` **—end—**

Example

The following example sets the following port attributes on port 2.1:

- disables auto-negotiation
- sets the description to 1234_West_Street

```
port set port 2.1 auto-neg off description
1234_West_Street
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 3-2
Resetting port attributes to default

Reset port attributes to default values.

Step	Action
1	Reset port attributes to default values: `port unset port <port> {description} {egress-mirror} {ingress-mirror} {mirror-encap} {mirror-encap-vid} {mirror-encap-tpid} {untagged-ctrl-vs} {untagged-data-vs} {untagged-data-vid} {mac-swap-vlan} {advertised-speed} {advertised-duplex}` where

port <port>	is the port.
description	is the port description.
egress-mirror	is egress port mirroring.
ingress-mirror	is ingress port mirroring.
mirror-encap	is port mirroring encapsulation.
mirror-encap-vid	is the port mirroring encapsulation VID.
mirror-encap-tpid	is the port mirroring encapsulation TPID.
untagged-ctrl-vs	is the virtual switch for untagged contol frames.
untagged-data-vs	is the virtual switch for untagged data frames.
untagged-data-vid	is the push/pop of VLAN ID for untagged data frames. ***Note:*** On the 3902, 3911, 3920, 3940 and 5140 platforms, the untagged-data-vid attribute pushes the VLAN ID on ingress but does not pop the VLAN ID on egress.
{mac-swap-vlan}	is the MAC SA/DA swap VLAN ID..
{advertised-speed}	is the advertised value for speed.
{advertised-duplex}	is the advertised value for duplex.

Example

The following example clears the descrption associated with port 1.

```
port unset port 1 description
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 3-3
Disabling a port

When you disable a port, the Link State administrative status is changed to disabled and the operational status shows disabled when the link is down.

When you disable a port directly, the transceiver is disabled.

Step	Action
1	Enter the following command: `port disable port <PortNameList>` where port <PortNameList> is the port to be disabled.

—end—

Example

The following example disables port 1.

```
port disable port 1
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 3-4
Enabling a port

When you enable a port, the Link State administrative status is changed to enabled and the operational status shows enabled when the link is up.

When you enable a port directly, the transceiver is enabled.

Step	Action
1	Enter the following command: `port enable port <port>` where port <port> is the port to be enabled.

—end—

Example

The following example enables port 1.

```
port enable port 1
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 3-5
Displaying port attributes

Display port attributes to

- verify configuration
- check link status
- troubleshoot issues related to the port

Step	Action

1 Display port attributes by entering the following command:

```
port show [port <port>]
```

where

port <port> is a 32-character string representing the name of the physical port or LAG.

—end—

Example

The following example shows sample output of the port show command.

```
> port show

+-------------------------------------------------------------------------------+
| Port Table      |           Operational Status            |  Admin Config     |
|-----------------+-----+--------------+----+---+--------+----+----+--------+----|
| Port    | Port  |     |  Link State  |    |   |        |Auto|    |        |Auto|
| Name    | Type  |Link |   Duration   |XCVR|STP| Mode   |Neg |Link| Mode   |Neg |
|---------+-------+-----+--------------+----+---+--------+----+----+--------+----|
| 1       |10/100/G|Down|   0d 0h 0m 0s|    |Dis|        |    |Ena |1000/FD | On |
| 2       | Gig   | Up  |   0d 3h56m14s|Ena |FWD|1000/FD | On |Ena |1000/FD | On |
...
+---------+-------+-----+--------------+----+---+--------+----+----+--------+----+
```

The following example shows sample output of the port show command for a specific port.

```
> port show port 1

+------------------------------ PORT 1 INFO -------------------------------+
| Field                     | Admin                  | Oper                 |
+---------------------------+------------------------+----------------------+
| Type                      | 10/100/G               |                      |
...
| Aggregation Membership    |                        |                      |
| VLAN Membership           | 1,10,100,1000          | 1,10,100,1000        |
+---------------------------+------------------------+----------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 3-6
Displaying port statistics

Two sets of statistics are stored:

- Current statistics, which are the values since the last statistics clear operation.
- Total statistics, which are the values since the last boot-up

The system also calculates throughput values to show current statistics in terms of rate.

You can display

- current statistics
- total statistics
- current throughput statistics
- statistics for specific ports

Step	Action

To display current statistics

1 Display current statistics:

```
port show statistics <statistics> [active] [delay
<NUMBER: 1-86400>] [count <NUMBER: 0-4294967295>] [scale
<tera|giga|mega|kilo|none>]
```

where

statistics <statistics>	is statistics collected since the last statistics clear operation.
active	displays active statistics.
delay<NUMBER: 1-86400	is the length of time in seconds to display the statistics.
count <NUMBER: 0-4294967295>	is the number of times to repeat the display.
scale <tera | giga | mega | kilo | none>	is the units for the dsplayed values.

Copyright© 2012-2013 Ciena® Corporation

To display total statistics

2 Display total statistics:

```
port show total-statistics <total-statistics> [active]
[delay <NUMBER: 1-86400>] [count <NUMBER: 0-4294967295>]
[scale <tera|giga|mega|kilo|none>]
```

where

total-statistics <total-statistics>	is all statistics collected since the last boot-up.
active	displays active statistics.
delay<NUMBER: 1-86400	is the length of time in seconds to display the statistics.
count <NUMBER: 0-4294967295>	is the number of times to repeat the display.
scale <tera | giga | mega | kilo | none>	is the units for the dsplayed values.

To display current throughput statistics

3 Display current throughput statistics:

```
port show throughput <throughput> [active] [delay
<NUMBER: 1-86400>] [count <NUMBER: 0-4294967295>] [scale
<tera|giga|mega|kilo|none>]
```

where

throughput <throughput>	is the port or ports to show.
active	displays active statistics.
delay<NUMBER: 1-86400	is the length of time in seconds to display the statistics.
count <NUMBER: 0-4294967295>	is the number of times to repeat the display.
scale <tera | giga | mega | kilo | none>	is the units for the dsplayed values. The default value is mega.

Copyright© 2012-2013 Ciena® Corporation

To display statistics for specific ports

4 Display statistics for specific ports:

```
port show port <port> [active] [count <NUMBER: 0-
4294967295>] [capabilities] [delay <NUMBER: 1-86400>]
[statistics] [total-statistics] [throughput] [scale
<tera|giga|mega|kilo|none>] [vlan] [type
<all|basic|errors|tx|rx>]
```

where

<port>	is the port or ports that you want to display port statistics for
active	displays active statistics.
count <NUMBER: 0-4294967295>	is the number of repetitions for throughput.
capabilities	displays port capabilities
delay<NUMBER: 1-86400>	is the time between samples in seconds.
statistics	displays port statistics.
total-statistics	displays total port statistics.
throughput	displays port throughput.
scale <tera | giga | mega | kilo | none>	is the scale used to show statistics.
vlan	displays port VLAN membership.
type <all | basic | errors | tx | rx>	is the type of statistics to show.

—end—

Examples

The following example shows sample output for all active port statistics summary.

```
> port show statistics active

+---------------------------- PORT STATISTICS SUMMARY ---------------------+
|   Port   |          Byte                 |          Pkt                  |
|          | Tx              | Rx          | Tx            | Rx            |
+---------+-----------------+-------------+---------------+---------------+
| 1        | 8326248088      | 0           | 67147162      | 0             |
| 2        | 8326247964      | 0           | 67147161      | 0             |
| 12       | 0               | 28879569152 | 0             | 225621634     |
+---------+-----------------+-------------+---------------+---------------+
```

The following example shows sample output for all active port total statistics summary.

```
> port show total-statistics active

+----------------------- PORT TOTAL STATISTICS SUMMARY --------------------
```


Copyright© 2012-2013 Ciena® Corporation

```
|   Port   |          Byte                  |       Pkt                  |
|          | Tx              | Rx           | Tx           | Rx          |
+----------+-----------------+--------------+--------------+-------------+
| 1        | 169713065596    | 1053155263708| 325536280    | 1372943892  |
| 2        | 21452132866     | 19878912     | 77950273     | 310608      |
| 12       | 885114437038    | 193572376072 | 1147664620   | 490131140   |
+----------+-----------------+--------------+--------------+-------------+
```

The following example shows sample output for all active port throughput statistics.

```
> port show throughput active

+--------------- PORT THROUGHPUT SUMMARY 5 SECOND SAMPLE -----------------+
|   Port   |          Bit Rate (Mbps)       |        Pkt Rate (Mpps)       |
|          | Tx              | Rx           | Tx              | Rx         |
+----------+-----------------+--------------+-----------------+------------+
| 1        |           0.867 |              |           0.001 |            |
| 2        |           0.867 |              |           0.001 |            |
| 12       |                 |        2.008 |                 |      0.003 |
+----------+-----------------+--------------+-----------------+------------+
```

The following example shows sample output for port 1 active statistics.

```
> port show port 1 statistics active

+--------------- PORT 1 STATISTICS -----+
| Statistic           | Value          |
+---------------------+----------------+
| RxBytes             | 9659942        |
| RxPkts              | 132603         |
| RxCrcErrorPkts      | 1              |
| RxUcastPkts         | 8499           |
| RxMcastPkts         | 81206          |
| RxBcastPkts         | 42897          |
| 64OctsPkts          | 127628         |
| 65To127OctsPkts     | 2345           |
| 256To511OctsPkts    | 2629           |
| 512To1023OctsPkts   | 1              |
| TxBytes             | 1485924        |
| TxPkts              | 12559          |
| TxUcastPkts         | 8233           |
| TxMcastPkts         | 4321           |
| TxBcastPkts         | 5              |
| Tx64OcPkts          | 4648           |
| Tx65To127OcPkts     | 3293           |
| Tx128To255OcPkts    | 4453           |
| Tx256To511OcPkts    | 51             |
| Tx512To1023OcPkts   | 56             |
| Tx1024To1518OcPkts  | 58             |
+---------------------+----------------+
```

The following example shows sample output for port 1 total active statistics.

```
> port show port 1 total-statistics active

+-------------------- PORT 1 STATISTICS --------------------+
| Statistic           | Total Value      | Value            |
+---------------------+------------------+------------------+
| RxBytes             | 9677152          | 9677152          |
| RxPkts              | 132855           | 132855           |
| RxCrcErrorPkts      | 1                | 1                |
| RxUcastPkts         | 8635             | 8635             |
```


Copyright© 2012-2013 Ciena® Corporation

```
| RxMcastPkts          | 81281             | 81281             |
| RxBcastPkts          | 42938             | 42938             |
| 64OctsPkts           | 127873            | 127873            |
| 65To127OctsPkts      | 2350              | 2350              |
| 256To511OctsPkts     | 2631              | 2631              |
| 512To1023OctsPkts    | 1                 | 1                 |
| TxBytes              | 1493626           | 1493626           |
| TxPkts               | 12654             | 12654             |
| TxUcastPkts          | 8328              | 8328              |
| TxMcastPkts          | 4321              | 4321              |
| TxBcastPkts          | 5                 | 5                 |
| Tx64OcPkts           | 4731              | 4731              |
| Tx65To127OcPkts      | 3300              | 3300              |
| Tx128To255OcPkts     | 4456              | 4456              |
| Tx256To511OcPkts     | 52                | 52                |
| Tx512To1023OcPkts    | 57                | 57                |
| Tx1024To1518OcPkts   | 58                | 58                |
+----------------------+-------------------+-------------------+
```

The following example shows sample output for port 1 active throughput statistics.

```
> port show port 1 throughput active

+---------------------- PORT 1 THROUGHPUT --------------------------------+
| Statistic            | Current Value    Delta Value      Rate Mpps & Mbps |
+----------------------+----------------+----------------+-----------------+
| Time                 |     1:19:56:19 |   0:00:14:20.0 |                 |
| RxBytes              |          9.702 |          0.072 |           0.000 |
| RxPkts               |          0.133 |          0.001 |           0.000 |
| RxUcastPkts          |          0.009 |          0.000 |           0.000 |
| RxMcastPkts          |          0.081 |          0.000 |           0.000 |
| RxBcastPkts          |          0.043 |          0.000 |           0.000 |
| 64OctsPkts           |          0.128 |          0.001 |           0.000 |
| 65To127OctsPkts      |          0.002 |          0.000 |           0.000 |
| 256To511OctsPkts     |          0.003 |          0.000 |           0.000 |
| TxBytes              |          1.500 |          0.027 |           0.000 |
| TxPkts               |          0.013 |          0.000 |           0.000 |
| TxUcastPkts          |          0.008 |          0.000 |           0.000 |
| Tx64OcPkts           |          0.005 |          0.000 |           0.000 |
| Tx65To127OcPkts      |          0.003 |          0.000 |           0.000 |
| Tx128To255OcPkts     |          0.004 |          0.000 |           0.000 |
| Tx256To511OcPkts     |          0.000 |          0.000 |           0.000 |
| Tx512To1023OcPkts    |          0.000 |          0.000 |           0.000 |
| Tx1024To1518OcPkts   |          0.000 |          0.000 |           0.000 |
+----------------------+----------------+----------------+-----------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 3-7
Monitoring port statistics

Use this procedure to continuously monitor port statistics for all ports or for specific ports. The system displays the statistics and automatically clears the screen before displaying the updated values.

To stop monitoring, press Ctrl+C.

Step	Action

To monitor all ports for current statistics

1 Monitor all ports for current statistics:

```
port monitor statistics <statistics> [active] [delay
<NUMBER: 1-86400>] [count <NUMBER: 0-4294967295>] [scale
<tera|giga|mega|kilo|none>]
```

where

statistics <statistics>	displays all statistics collected from the last statistics clear operation.
active	displays active statistics.
delay<NUMBER: 1-86400	is the length of time in seconds to display the statistics.
count <NUMBER: 0-4294967295>	is the number of times to repeat the display.
scale <tera | giga | mega | kilo | none>	is the units for the dsplayed values.

To monitor all ports for total statistics

2 Monitor all ports for total statistics:

```
port monitor total-statistics <total-statistics> [active]
[delay <NUMBER: 1-86400>] [count <NUMBER: 0-4294967295>]
[scale <tera|giga|mega|kilo|none>]
```

where

total-statistics <total-statistics>	displays all statistics collected since the last boot-up.
active	displays active statistics.

Copyright© 2012-2013 Ciena® Corporation

where

delay<NUMBER: 1-86400	is the length of time in seconds to display the statistics.
count <NUMBER: 0-4294967295>	is the number of times to repeat the display.
scale <tera | giga | mega | kilo | none>	is the units for the dsplayed values.

To monitor all ports for throughput statistics

3 Monitor all ports for throughput statistics:

```
port monitor throughput <throughput> [active] [delay
<NUMBER: 1-86400>] [count <NUMBER: 0-4294967295>] [scale
<tera|giga|mega|kilo|none>]
```

where

throughput <throughput>	displays current throughput statistics.
active	displays active statistics.
delay<NUMBER: 1-86400	is the length of time in seconds to display the statistics.
count <NUMBER: 0-4294967295>	is the number of times to repeat the display.
scale <tera | giga | mega | kilo | none>	is the units for the dsplayed values.

To monitor specific ports

4 Monitor specific ports:

```
port monitor port <port> [active] [delay <NUMBER: 1-
86400>] [scale <tera|giga|mega|kilo|none>] {statistics]
[total-statistics] [throughput] [type
<all|basic|errors|tx|rx>]
```

where

port <port>	is the port or ports that you want to monitor
active	displays active statistics.
delay<NUMBER: 1-86400	is the length of time in seconds to display the statistics.
scale <tera | giga | mega | kilo | none>	is the units for the dsplayed values.
statistics	displays port statistics.

Copyright© 2012-2013 Ciena® Corporation

where

total-statistics	displays total port statistics.
throughput	displays port throughput.
type <all | basic | errors | tx | rx>	is the type of statistics to show.

—end—

Examples

The following example shows sample output of monitoring total statistics for all active ports.

```
> port monitor total-statistics active delay 10
<Screen clears>

+---------------------- PORT TOTAL STATISTICS SUMMARY ---------------------+
|   Port   |          Byte                    |          Pkt                |
|          | Tx              | Rx             | Tx            | Rx          |
+----------+-----------------+----------------+---------------+-------------+
| 1        | 170475797780    | 1053155263708  | 331687346     | 1372943892  |
| 2        | 22214864802     | 19878912       | 84101337      | 310608      |
| 12       | 885114437038    | 196217923208   | 1147664620    | 510799477   |
+----------+-----------------+----------------+---------------+-------------+
```

The following example shows sample output of monitoring statistics for all ports.

```
> port monitor statistics active delay 10
<Screen clears>

+--------------------------- PORT STATISTICS SUMMARY ----------------------+
|   Port   |          Byte                    |          Pkt                |
|          | Tx              | Rx             | Tx            | Rx          |
+----------+-----------------+----------------+---------------+-------------+
| 2        | 36938335        | 67861410       | 224689        | 1023888     |
| 3        | 3522850         | 21965574       | 22329         | 228469      |
+----------+-----------------+----------------+---------------+-------------+
```

The following example shows sample output of monitoring throughput statistics for all ports.

```
> port monitor throughput active
Info: This CLI output may take a while to display press CTRL-C to abort

<Screen clears>
+------------- PORT THROUGHPUT SUMMARY   5 SECOND SAMPLE -----------------+
|   Port   |          Bit Rate (Mbps)         |       Pkt Rate (Mpps)       |
|          | Tx              | Rx             | Tx            | Rx          |
+----------+-----------------+----------------+---------------+-------------+
| 1        |           0.860 |                |         0.001 |             |
| 2        |           0.860 |                |         0.001 |             |
| 12       |                 |          2.984 |               |       0.003 |
+----------+-----------------+----------------+---------------+-------------+
```

The following example shows sample output of monitoring total statistics for a specific port.

```
> port monitor port 1 total-statistics active delay 10
```


Copyright© 2012-2013 Ciena® Corporation

```
<Screen clears>

INFO: Waiting 10 seconds for display. Abort with CTRL-c

+--------------------- PORT 1 STATISTICS --------------------+
| Statistic           | Total Value       | Value             |
+---------------------+-------------------+-------------------+
| RxBytes             | 9844868           | 9844868           |
| RxPkts              | 135233            | 135233            |
| RxCrcErrorPkts      | 1                 | 1                 |
| RxUcastPkts         | 9572              | 9572              |
| RxMcastPkts         | 82222             | 82222             |
| RxBcastPkts         | 43438             | 43438             |
| 64OctsPkts          | 130164            | 130164            |
| 65To127OctsPkts     | 2406              | 2406              |
| 256To511OctsPkts    | 2662              | 2662              |
| 512To1023OctsPkts   | 1                 | 1                 |
| TxBytes             | 1580957           | 1580957           |
| TxPkts              | 13396             | 13396             |
| TxUcastPkts         | 9070              | 9070              |
| TxMcastPkts         | 4321              | 4321              |
| TxBcastPkts         | 5                 | 5                 |
| Tx64OcPkts          | 5195              | 5195              |
| Tx65To127OcPkts     | 3519              | 3519              |
| Tx128To255OcPkts    | 4480              | 4480              |
| Tx256To511OcPkts    | 65                | 65                |
| Tx512To1023OcPkts   | 60                | 60                |
| Tx1024To1518OcPkts  | 77                | 77                |
+---------------------+-------------------+-------------------+
```


Copyright© 2012-2013 Ciena® Corporation

The following example shows sample output for monitoring statistics for a specific port.

```
> port monitor port 1 statistics active delay 10
<Screen clears>

INFO: Waiting 10 seconds for display. Abort with CTRL-c

+--------------- PORT 1 STATISTICS -----+
| Statistic           | Value           |
+---------------------+-----------------+
| RxBytes             | 9866442         |
| RxPkts              | 135531          |
| RxCrcErrorPkts      | 1               |
| RxUcastPkts         | 9648            |
| RxMcastPkts         | 82368           |
| RxBcastPkts         | 43514           |
| 64OctsPkts          | 130448          |
| 65To127OctsPkts     | 2415            |
| 256To511OctsPkts    | 2667            |
| 512To1023OctsPkts   | 1               |
| TxBytes             | 1586934         |
| TxPkts              | 13454           |
| TxUcastPkts         | 9128            |
| TxMcastPkts         | 4321            |
| TxBcastPkts         | 5               |
| Tx64OcPkts          | 5238            |
| Tx65To127OcPkts     | 3530            |
| Tx128To255OcPkts    | 4482            |
| Tx256To511OcPkts    | 66              |
| Tx512To1023OcPkts   | 60              |
| Tx1024To1518OcPkts  | 78              |
+---------------------+-----------------+
```

The following example shows sample output for monitoring throughput statistics for a specific port.

```
> port monitor port 1 throughput active
<Screen clears>

INFO: Waiting 5 seconds for display. Abort with CTRL-c

+---------------------- PORT 1 THROUGHPUT ---------------------------------+
| Statistic         | Current Value  | Delta Value      | Rate Mpps & Mbps |
+-------------------+----------------+------------------+------------------+
| Time              |     1:20:30:17 |     0:00:00:05.0 |                  |
| RxBytes           |          9.887 |            0.000 |            0.000 |
| RxPkts            |          0.136 |            0.000 |            0.000 |
| RxUcastPkts       |          0.010 |            0.000 |            0.000 |
| RxMcastPkts       |          0.083 |            0.000 |            0.000 |
| RxBcastPkts       |          0.044 |            0.000 |            0.000 |
| 64OctsPkts        |          0.131 |            0.000 |            0.000 |
| TxBytes           |          1.593 |            0.000 |            0.000 |
| TxPkts            |          0.014 |            0.000 |            0.000 |
| TxUcastPkts       |          0.009 |            0.000 |            0.000 |
| Tx65To127OcPkts   |          0.004 |            0.000 |            0.000 |
+-------------------+----------------+------------------+------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 3-8
Clearing current statistics

Clear current statisitcs when you no longer want to view them. Clearing current statistics does not clear total statistics.

You can clear current statisitcs for

- all ports
- specific ports

Step	Action

To clear current statistics for all ports

1 Clear current statistics for all ports:

```
port clear statistics
```

To clear current statistics for specific ports

2 Clear current statistics for specific ports:

```
port clear port <port> statistics
```

where

port <port> is the port or ports that you want to clear statistics for.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 3-9
Displaying blade information

Display blade information.

Step	Action

1 Display blade information:

```
blade show [attributes] [capabilities] [information]
[state]
```

where

attributes	displays hardware device identification information. Hardware device identification information is device type, hardware version, serial number, MAC address, manufactured date, and E-PROM (param) version. The blade show attributes output varies depending on the hardware platform. The CLEI Code is only displayed for 3916, 3930, 3931, and 5150 devices
capabilities	displays the capabilities of the board and ports, including blade type, RAM and flash file sizes, port types supported, and enhanced ports.
information	displays general blade information, including blade type, number of ports, MAC address, administrative and operational states, and date of last reboot.
state	displays administrative and operational states.

—end—

Examples

The following example shows sample output for a single blade device.

```
> blade show

+----------------------- BLADE SUMMARY -----------------------+
| Slot | Ports | PortBaseMac       | BladeType | OperState   |
|------+-------+-------------------+-----------+-------------|
| 1    | 12    | 00:03:18:55:71:d2 | Single    | Enabled     |
+------+-------+-------------------+-----------+-------------+
```

The following example shows all extended blade details in one command (on a 5150).

```
> blade show attributes capabilities information state

//The blade attributes:

+----------------- BLADE DEVICE ID -----------------+
| Parameter                 | Value                 |
```


Copyright© 2012-2013 Ciena® Corporation

```
+--------------------------+-----------------------+
| Board Device Type        | 091                   |
| Board Hardware Version   | 1705150830/004        |
| Board Serial Number      | B6054200              |
| Board MAC Address        | 00:03:18:ac:e9:40     |
| Manufactured Date        | 11-11-2010            |
| CLEI Code                | COMP100BRA            |
| Location of Manufacture  | 1                     |
| Module Part Num          | 1705150900/009        |
| Module Serial Num        | M6145938              |
| Param Version            | 007                   |
+--------------------------+-----------------------+

+--------------- MODULE 2 DEVICE ID ----------------+
| Parameter                | Value                 |
+--------------------------+-----------------------+
| Board Device Type        | 092                   |
| Board Hardware Version   | 1705100810/006        |
| Board Serial Number      | B6061295              |
| Manufactured Date        | 22-10-2010            |
| CLEI Code                | COUIA3CPAA            |
| Location of Manufacture  | 1                     |
| Module Part Num          | 1705100900/004        |
| Module Serial Num        | M6154551              |
| Param Version            | 007                   |
+--------------------------+-----------------------+

+--------------- MODULE 3 DEVICE ID ----------------+
| Parameter                | Value                 |
+--------------------------+-----------------------+
| Board Device Type        | 093                   |
| Board Hardware Version   | 1705101810/003        |
| Board Serial Number      | B6004593              |
| Manufactured Date        | 04-07-2010            |
| CLEI Code                | COUIA7APAA            |
| Location of Manufacture  | 1                     |
| Module Part Num          | 1705101900/002        |
| Module Serial Num        | M6043072              |
| Param Version            | 007                   |
+--------------------------+-----------------------+

//The blade capabilities:

+---------------------------- BOARD CAPABILITIES -----------------------------+
| Parameter          | Value                                                  |
+--------------------+--------------------------------------------------------+
| Capability Class   | 0                                                      |
| Board Type         | 091                                                    |
| Board Name         | CN 5150                                                |
| Board Description  | CN 5150 Service Aggregation Switch                     |
| Blade Type         | Single                                                 |
| No. Ports          | 52                                                     |
| Has Dcard          | No                                                     |
| Address Ram Size   | 0x0                                                    |
| Boot Flash Size    | 0x200000                                               |
| Packet Ram Size    | 0x0                                                    |
| Program Ram Size   | 0x2000000                                              |
+--------------------+--------------------------------------------------------+
| No. 10 Gig Ports   | 4                                                      |
| No. Gig Ports      | 48                                                     |
| No. Fe Ports       | 0                                                      |
| No. 100Fx Ports    | 0                                                      |
| No. Eth Ports      | 0                                                      |
| Total Ports        | 52                                                     |
| Enhanced Ports List| 2.1 2.2                                                |
+--------------------+--------------------------------------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

```
+------------------------- SUMMARY PORT CAPABILITIES ---------------------------------+
| Port| Type         | Speed                    | Duplex   | Aneg  | Pause             | Enh |
+-----+--------------+--------------------------+----------+-------+-------------------+---+
| 1.1 | 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.2 | 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.3 | 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.4 | 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.5 | 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.6 | 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.7 | 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.8 | 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.9 | 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.10| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.11| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.12| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.13| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.14| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.15| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.16| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.17| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.18| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.19| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.20| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.21| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.22| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.23| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.24| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.25| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.26| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.27| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.28| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.29| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.30| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.31| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.32| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.33| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.34| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.35| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.36| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.37| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.38| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.39| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.40| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.41| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.42| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.43| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.44| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.45| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.46| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.47| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 1.48| 100/G        | 100Mbps,1Gig,Auto        | half,full| on,off| off,sym,a-rx      |No |
| 2.1 | 10GigEthernet| 10Gig                    | full     | N/A   | off               |Yes|
| 2.2 | 10GigEthernet| 10Gig                    | full     | N/A   | off               |Yes|
| 3.1 | 10GigEthernet| 10Gig                    | full     | N/A   | off               |No |
| 3.2 | 10GigEthernet| 10Gig                    | full     | N/A   | off               |No |
+-----+--------------+--------------------------+----------+-------+-------------------+---+

//The blade information:

+------------------ BLADE INFO --------------------+
| Parameter              | Value                   |
+------------------------+-------------------------+
| Slot                   | 1                       |
| Blade Type             | Single                  |
| Number of Ports        | 52                      |
| Port Base MAC Address  | 00:03:18:ac:e9:42       |
| Admin State            | Enabled                 |
| Oper State             | Enabled                 |
```


Copyright© 2012-2013 Ciena® Corporation

```
| Last Reboot           | Thu Jan  1 00:00:00 1970 |
+-----------------------+--------------------------+

//The blade state:

+----- BLADE STATE ------+
| AdminState | OperState |
+------------+-----------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 3-10
Displaying port capabilities

You can display:

- port capabilities for the chassis and a summary of port capabilities
- capabilities for a specified port

Step	Action

To display port capabilities for the chassis and a summary of port capabilities

1 Display chassis port capabilities and a summary of all port capabilities:

```
port show capabilities
```

To display capabilities for a specific port

2 Display capabilities for a specific port:

```
port show port <port> capabilities
```

where

port <port> is a 32-character string representing the name of the physical port or LAG.

—end—

Examples

The following example shows sample output for the port show capabilities command.

```
> port show capabilities

+--------------------+-----------------------------------------------------+
| No. 10 Gig Ports   | 4                                                   |
| No. Gig Ports      | 8                                                   |
| No. Fe Ports       | 0                                                   |
| No. 100Fx Ports    | 0                                                   |
| No. Eth Ports      | 0                                                   |
| Total Ports        | 12                                                  |
+--------------------+-----------------------------------------------------+

+------------------------- SUMMARY PORT CAPABILITIES -------------------------+
| Port| Type          | Speed                | Duplex    | Aneg   | Pause          |
+-----+---------------+----------------------+-----------+--------+----------------+
| 1   | 10/100/G      | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx    |
| 2   | 10/100/G      | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx    |
| 3   | 10/100/G      | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx    |
| 4   | 10/100/G      | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx    |
| 5   | 10/100/G      | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx    |
| 6   | 10/100/G      | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx    |
| 7   | 10/100/G      | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx    |
| 8   | 10/100/G      | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx    |
| 9   | 10GigEthernet | 10Gig                | full      | N/A    | off             |
| 10  | 10GigEthernet | 10Gig                | full      | N/A    | off             |
| 11  | 10GigEthernet | 10Gig                | full      | N/A    | off             |
```


Copyright© 2012-2013 Ciena® Corporation

```
| 12  | 10GigEthernet| 10Gig               | full     | N/A   | off            |
+-----+--------------+---------------------+----------+-------+----------------+
```

The following example shows sample output for the port show capabilities command applied to a specified port.

```
> port show port 1 capabilities

+------------------------ PORT 1 CAPABILITIES --------------------+
| Field                          | Value                          |
+--------------------------------+--------------------------------+
| Port Number                    | 1                              |
| Port Type                      | 10/100/G                       |
| Port Speed                     | 10Mbps,100Mbps,1Gig,Auto       |
| Port Duplex                    | half,full                      |
| Port Auto Negotiation          | on,off                         |
| Port Pause Advertisement       | off,sym,a-tx,s-a-rx            |
| Port Pause                     | off,sym,a-rx                   |
| Port Feature Capabilities      | Normal                         |
+--------------------------------+--------------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 3-11
Displaying port Ethernet configuration

You can display port attributes for Ethernet configuration for all or for a specific line module, including name, type, admin status, speed, duplex, flow control, flow control advertised, auto negotiation, and MTU size.

Step	Action
1	Display port attritbutes for Ethernet configuration: `port show ethernet-config`

—end—

Example

The following example shows sample output for the port show ethernet-config command.

```
> port show ethernet-config
+------------------------ PORT ETHERNET CONFIGURATION -------------------------+
|          |          |          |         |       |       |      |      |      Mirror      |
|  Port    |  Port    |  Admin   |         |       |       |  FC  | Auto |  MTU |  Status  |
|  Name    |  Type    |  Status  |  Speed  |  Dplx |  FC   |  Adv | Neg  |  Size|State|Eg|Ig|
+----------+----------+----------+---------+-------+-------+------+------+------+-----+--+--+
|  1       |  Gig     |  Ena     |  1000   |   ??  |  off  |  off |  On  | 1526 | Off | 0| 0|
|  2       |  Gig     |  Ena     |  1000   |  Full |  off  |  off |  On  | 1526 | Off | 0| 0|
|  3       |  Gig     |  Ena     |  1000   |  Full |  off  |  off |  On  | 1526 | Off | 0| 0|
|  4       | 10/100/G |  Ena     |  1000   |   ??  |  off  |  off |  On  | 1526 | Off | 0| 0|
|  5       | 10/100/G |  Ena     |  1000   |   ??  |  off  |  off |  On  | 1526 | Off | 0| 0|
|  6       | 10/100/G |  Ena     |  1000   |   ??  |  off  |  off |  On  | 1526 | Off | 0| 0|
|  7       | 10/100/G |  Ena     |  1000   |   ??  |  off  |  off |  On  | 1526 | Off | 0| 0|
|  8       | 10/100/G |  Ena     |  1000   |   ??  |  off  |  off |  On  | 1526 | Off | 0| 0|
|  9       |  10Gig   |  Ena     |    10G  |   ??  |  off  |  off |  Off | 1526 | Off | 0| 0|
|  10      |  10Gig   |  Ena     |    10G  |   ??  |  off  |  off |  Off | 1526 | Off | 0| 0|
|  11      |  10Gig   |  Ena     |    10G  |   ??  |  off  |  off |  Off | 1526 | Off | 0| 0|
|  12      |  10Gig   |  Ena     |    10G  |   ??  |  off  |  off |  Off | 1526 | Off | 0| 0|
+----------+----------+----------+---------+-------+-------+------+------+------+-----+--+--+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 3-12
Displaying port status

Port status included the operational information, such as the link state, link state duration, whether transceivers are enabled or disabled, speed, duplex, maximum frame size, and flow control. You can display the status for all ports or ports on a specific line module.

Step	Action
1	Display port status: `port show status`

—end—

Example

The following example shows sample output for the port show status command.

```
> port show status
+------------------------ PORT OPERATIONAL STATUS -------------------------+
|  |                            |      |Link State|     |     |Speed/   |MTU  |Flow  |
|##| Description                |Link  |Duration  |XCVR |STP  |Duplex   |Size |Ctrl  |
+--+----------------------------+------+----------+-----+-----+---------+-----+------+
|1 |                            |Down  |  1d21h35m|     |Dis  |         |1526 |      |
|2 |                            | Up   |  1d22h37m|Ena  |FWD  |1000/FD  |1526 |off   |
|3 |                            | Up   |  1d21h33m|Ena  |FWD  |1000/FD  |1526 |off   |
|4 |                            |Down  |  0h 0m 0s|     |Dis  |         |1526 |      |
|5 |                            |Down  |  0h 0m 0s|     |Dis  |         |1526 |      |
|6 |                            |Down  |  0h 0m 0s|     |Dis  |         |1526 |      |
|7 |                            |Down  |  0h 0m 0s|     |Dis  |         |1526 |      |
|8 |                            |Down  |  0h 0m 0s|     |Dis  |         |1526 |      |
|9 |                            |Down  |  0h 0m 0s|     |Dis  |         |1526 |      |
|10|                            |Down  |  0h 0m 0s|     |Dis  |         |1526 |      |
|11|                            |Down  |  0h 0m 0s|     |Dis  |         |1526 |      |
|12|                            |Down  |  0h 0m 0s|     |Dis  |         |1526 |      |
+--+----------------------------+------+----------+-----+-----+---------+-----+------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 3-13
Displaying a list of supported optics

Small Form-factor Pluggable (SFP) and 10 Gigabit SFPs (XFP) devices are hot-swappable compact optical transceivers. Port transceiver information, including status and type, is available via CLI or SNMP.

The system software supports transceivers that are compliant with the following documents:

- XFP Xcvr spec SFF INF 8077i Rev 4.5, Tunable Xcvr spec SFF-8477 Rev 1.3 Draft.
- Small Form Factor Pluggable (SFP) Transceiver Multi Source Agreement, September 14, 2000
- Digital Diagnostic Monitoring Interface for Optical transceivers SFF-8473, Draft Revision 9.0, April 4, 2002.

Step	Action
1	Display a list of supported optics: `port xcvr show supported`

—end—

Example

The following example shows sample output for the port xcvr show supported command.

```
> port xcvr show supported

+----All Supported Transceivers----+
|                                   |
|     Part Number      | XCVR Speed |
+--------------------+-------------+
|        XCVR-010X31   |     100M    |
|        XCVR-040X31   |     100M    |
|        XCVR-040R55   |     100M    |
|        XCVR-040R31   |     100M    |
|        XCVR-010S55   |     100M    |
|        XCVR-010S31   |     100M    |
|        XCVR-010L31   |     100M    |
|        XCVR-040L31   |     100M    |
|        XCVR-100D43   |   100M / 1G |
|        XCVR-100D45   |   100M / 1G |
|        XCVR-100D47   |   100M / 1G |
...                                 |
+--------------------+-------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 3-14
Displaying transceiver information

Diagnostics can be displayed for transceivers that support diagnostics. If the transceiver does not have internal diagnostics capabilities, an error is returned.

You can display

- diagnostic information for a specified port
- a summary of transceiver status
- a summary of transceiver status for a specific port
- vendor EPROM data for a specific port

Step	Action

To display diagnostic information for a specified port

1 Display diagnostic information for a specified port:

```
port show port <port> diagnostics
```

where

port <port> is the port or ports to show.

To display a summary of transceiver status

2 Display a summary of transceiver status:

```
port xcvr show
```

To display a summary of transceiver status for a specific port

3 Display a summary of transceiver status for a specific port:

```
port xcvr show port <port> state
```

where

port <port> is the port or ports to show.

To display vendor EPROM data for a specific port

4 Display vendor EPROM data for a specific port:

```
port show port <port> vendor
```

where

port <port> is the port or ports to show.

—end—

Copyright© 2012-2013 Ciena® Corporation

Example

The following example shows sample output for a transceiver that supports diagnostics.

```
> port xcvr show port 3 diagnostics

+--------------------- XCVR DIAGNOSTICS - Port 3        -----------------+
|                    |           |          Alarm       |        Warning       |
| Output             | Value     | Threshold      | Flag | Threshold      | Flag |
+--------------------+-----------+----------------+------+----------------+------+
| Temp     (degC)    |  42.031   | HIGH 105.000   | 0    | HIGH 100.000   | 0    |
|                    |           | LOW  -45.000   | 0    | LOW  -40.000   | 0    |
+--------------------+-----------+----------------+------+----------------+------+
| Vcc      (volts)   |   3.299   | HIGH   3.630   | 0    | HIGH   3.460   | 0    |
|                    |           | LOW    2.970   | 0    | LOW    3.130   | 0    |
+--------------------+-----------+----------------+------+----------------+------+
| Bias       (mA)    |   4.112   | HIGH  12.000   | 0    | HIGH  10.000   | 0    |
|                    |           | LOW    1.000   | 0    | LOW    2.000   | 0    |
+--------------------+-----------+----------------+------+----------------+------+
| Tx Power   (mW)    |   0.257   | HIGH   1.412   | 0    | HIGH   0.707   | 0    |
|                    |           | LOW    0.056   | 0    | LOW    0.112   | 0    |
+--------------------+-----------+----------------+------+----------------+------+
| Tx Power (dBm)     |  -5.8905  | HIGH  +1.4999  | 0    | HIGH  -1.5003  | 0    |
|                    |           | LOW  -12.5026  | 0    | LOW   -9.5001  | 0    |
+--------------------+-----------+----------------+------+----------------+------+
| Rx Power   (mW)    |   0.0574  | HIGH   1.9954  | 0    | HIGH   1.0000  | 0    |
|                    |           | LOW    0.0100  | 0    | LOW    0.0200  | 0    |
+--------------------+-----------+----------------+------+----------------+------+
| Rx Power (dBm)     | -12.4109  | HIGH  +3.0001  | 0    | HIGH  +0.0000  | 0    |
|                    |           | LOW  -20.0000  | 0    | LOW  -16.9897  | 0    |
+--------------------+-----------+----------------+------+----------------+------+
```

The following example shows sample output for a summary of transceiver status.

```
> port xcvr show

+----+-----+-----+---------Transceiver-Status------------+-----------------+----+
|    |Admin| Oper|                                       |Ether Medium &   |Diag|
|Port|State|State|      Vendor Name & Part Number        |Connector Type   |Data|
+----+-----+-----+---------------------------------------+-----------------+----+
|1   |Empty|     |                                       |                 |    |
|2   |Empty|     |                                       |                 |    |
|3   |Ena  |Ena  |CIENA-FBX XCVR-A00G85 Rev10            |1000BASE-SX/LC   |Yes |
|4   |Ena  |Ena  |CIENA-FBX XCVR-A00G85 Rev10            |1000BASE-SX/LC   |Yes |
|5   |Ena  |Ena  |CIENA-FBX XCVR-A00G85 Rev10            |1000BASE-SX/LC   |Yes |
|6   |Ena  |Ena  |CIENA-LMT XCVR-A80D43 RevA             |1000BASE-LX/LC   |Yes |
+----+-----+-----+---------------------------------------+-----------------+----+
```

The following example shows sample output for a summary of transceiver status for a specific port.

```
> port xcvr show port 12 state

+----+-----+-----+---------Transceiver-Status--------+---------------+----+
|    |Admin| Oper|                                   |Ether Medium & |Diag|
|Port|State|State|      Vendor Name & Part Number    |Connector Type |Data|
+----+-----+-----+-----------------------------------+---------------+----+
|12  |Ena  |Ena  |CIENA XCVR-010Y31 Rev10            |1000BASE-LX/LC |    |
+----+-----+-----+-----------------------------------+---------------+----+
```


Copyright© 2012-2013 Ciena® Corporation

The following example shows sample output for vendor EPROM data for a specific port.

```
> port xcvr show port 7 vendor

+------------------------ XCVR VENDOR DATA - Port 7    -------------------------+
| Parameter                | Value              | Decoded String Equivalent    |
+--------------------------+--------------------+------------------------------+
| Identifier               | 0x3                | SFP transceiver              |
| Ext. Identifier          | 0x4                | SFP/GBIC                     |
| Connector                | 0x7                | LC                           |
+--------------------------+--------------------+------------------------------+
| Transceiver Codes        | 0x010d001202000000 |                              |
|  - 10 GbE Compliance     | 0x00               |                              |
|  - SONET Compliance      | 0x0000             |                              |
|  - Ethernet Compliance   | 0x02               | 1000BASE-LX                  |
|  - Link Length           | 0x12               | Long distance (L)            |
|  - Transmitter Technology| 0x0012             | Longwave laser (LL)          |
|  - Transmission Media    | 0x0d               | Single Mode (SM)             |
|  - Channel speed         | 0x01               | 100 MBytes/Sec               |
+--------------------------+--------------------+------------------------------+
| Encoding                 | 0x01               | 8B10B                        |
| BR, Nominal              | 13                 | Gigabit                      |
|--------------------------+--------------------+------------------------------|
| Length(9um fiber) 1km    | 10                 | 10km                         |
| Length(9um fiber) 100m   | 100                | 10000m                       |
| Length(50um) 10m         | 55                 | 550m                         |
| Length(62.5um) 10m       | 55                 | 550m                         |
| Length(copper) 1m        | 0                  | 0m                           |
|--------------------------+--------------------+------------------------------|
| Vendor Name              | CIENA              |                              |
| Vendor OUI               | 0x000000           |                              |
| Vendor PN                | XCVR-010M31-03     |                              |
| Vendor Revision          | 10                 |                              |
| Wavelength               | 0                  |                              |
|--------------------------+--------------------+------------------------------|
| Options                  | 0x1a               |                              |
|  - RATE_SELECT           | Bit 5              | No                           |
|  - TX_DISABLE            | Bit 4              | Yes                          |
|  - TX_FAULT              | Bit 3              | Yes                          |
|  - Loss of Signal Invert | Bit 2              | No                           |
|  - Loss of Signal        | Bit 1              | Yes                          |
|--------------------------+--------------------+------------------------------|
| BR, max                  | 0                  |                              |
| BR, min                  | 0                  |                              |
| Vendor Serial Number     | A9640060800547     |                              |
| Date (mm/dd/yy)          | 03/15/06           |                              |
|--------------------------+--------------------+------------------------------|
| Diag Monitor Type        | 0x0                |                              |
|  - Legacy diagnostics    | Bit 7              | No                           |
|  - Diagnostics monitoring| Bit 6              | No                           |
|  - Internally calibrated | Bit 5              | No                           |
|  - Externally calibrated | Bit 4              | No                           |
|  - Rx power measurement  | Bit 3              | OAM                          |
|--------------------------+--------------------+------------------------------|
| Enhanced Options         | 0x0                |                              |
|  - Alarm/Warning Flags   | Bit 7              | No                           |
|  - Soft TX_DISABLE       | Bit 6              | No                           |
|  - Soft TX_FAULT         | Bit 5              | No                           |
|  - Soft RX_LOS           | Bit 4              | No                           |
|  - Soft RATE_SELECT      | Bit 3              | No                           |
|--------------------------+--------------------+------------------------------|
| SFF-8472 Compliance      | 0x0                | None                         |
+--------------------------+--------------------+------------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 3-15
Determining transceiver speed

When a transceiver is plugged in, the port speed is blank until a link is established, and then it is set to match the transceiver speed.

The Encoding “Value” column displays the actual value read from the optic, while the “Decoded String Equivalent” column indicates the supported port speed.

Step	Action
1	Display transceiver information: `port xcvr show`
2	Display transceiver vendor data: `port xcvr show port <port> vendor`

where

port <port>	is the port.
vendor	displays transceiver vendor data.

—end—

Examples

The following example shows sample output for the port show command.

```
> port show
```

Port Table		Operational Status					Admin Config		
Port Name	Port Type	Link	Link State Duration	XCVR	STP	Mode	Auto Neg	Link	Mode	Auto Neg
1	10/100/G	Up	0d 5h 4m27s		FWD	100/FD	On	Ena	1000/FD	On
2	10/100/G	Up	0d 5h 4m27s		FWD	100/FD	On	Ena	1000/FD	On
3	10/100/G	Down	0d 0h 0m 0s		Dis		On	Ena	1000/FD	On
4	10/100/G	Down	0d 0h 0m 0s		Dis		On	Ena	1000/FD	On
5	10/100/G	Down	0d 0h 0m 0s		Dis		On	Ena	1000/FD	On
6	10/100/G	Down	0d 0h 0m 0s		Dis		On	Ena	1000/FD	On
7	10/100/G	Down	0d 0h 0m 0s		Dis		On	Ena	1000/FD	On
8	10/100/G	Down	0d 0h 0m 0s		Dis		On	Ena	1000/FD	On
9	10/100/G	Down	0d 0h 0m 0s		Dis		On	Ena	Auto/FD	On
10	10/100/G	Down	0d 0h 0m 0s		Dis		On	Ena	Auto/FD	On
11	10/100/G	Down	0d 0h 0m 0s		Dis		On	Ena	Auto/FD	On
12	Gig	Up	4d 2h48m52s	Ena	FWD	**1000/FD**	On	Ena	Auto/FD	On

The following example shows sample output for the port xcvr show port <PortNameList> vendor command.

Copyright© 2012-2013 Ciena® Corporation

```
> port xcvr show port 1 vendor

+------------------------ XCVR VENDOR DATA - Port 1    -------------------------+
| Parameter                | Value               | Decoded String Equivalent    |
+--------------------------+---------------------+------------------------------+
| Identifier               | 0x3                 | SFP transciever              |
| Ext. Identifier          | 0x4                 | SFP/GBIC                     |
| Connector                | 0x7                 | LC                           |
+--------------------------+---------------------+------------------------------+
| Transceiver Codes        | 0x0000000002000000  |                              |
|  - SONET Compliance      | 0x0000              |                              |
|  - Ethernet Compliance   | 0x02                | 1000BASE-LX                  |
|  - Link Length           | 0x00                | unknown                      |
|  - Transmitter Technology| 0x0000              | unknown                      |
|  - Transmission Media    | 0x00                | unknown                      |
|  - Channel speed         | 0x00                | unknown                      |
+--------------------------+---------------------+------------------------------+
| Encoding                 | 0x01                | 8B10B                        |
| BR, Nominal              | 13                  | Gigabit                      |
|--------------------------+---------------------+------------------------------|
...
+--------------------------+---------------------+------------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 3-16
Tuning XFP transceivers

Tunable XFPs provide the ability to set the laser frequency in gigahertz (GHz), wavelength in nanometers with decimals, or set the channel. When you set the value for the frequency, wavelength, or channel, the other parameters are automatically populated.

If the range is set out of the supported range for the XFP, an error message is returned with the correct range. Attempting to set the tuning parameters for an SFP that does not support tunability generates an error.

Step	Action

To tune XFP transceivers

1 View the current values for frequency, wavelength or channel:

```
port xcvr show port <port> tunability
```

where

port <port>	is the port.
tunability	displays transceiver tunable data.

2 Set the frequency, wavelength, or channel :

```
port xcvr set port <port> {[frequency <NUMBER>] |
[wavelength <String>] | [channel <NUMBER>]}
```

where

port <port>	is the port or ports to set.
frequency <NUMBER>	is the transceiver frequency in GHz.
wavelength <String>	is the transceiver wavelength in nanometers.
channel <NUMBER>	is the transceiver channel number.

Copyright© 2012-2013 Ciena® Corporation

To unset transceiver values

3 Unset transceiver values:

```
port xcvr unset port <port> {frequency} {wavelength}
{channel}
```

where

port <port>	is the port or ports to unset.
frequency <NUMBER>	is the transceiver frequency.
wavelength <String>	is the transceiver wavelength.
channel <NUMBER>	is the transceiver channel number.

—end—

Examples

The folowing example shows sample output for a port with a tunable XFP.

```
> port xcvr show port 2.1 tunability

+--------------------- XCVR Tunability - Port 2.1      -----------------+
| Field                  |                    Value                     |
+------------------------+----------------------------------------------+
| Frequency Tunable      |                      Yes                     |
+------------------------+-----------------+-----------------+----------+
|                        |          GHz    |           nm    |   ch#    |
|                        +-----------------+-----------------+----------+
| Admin                  |       191100    |       1568.8    |      1   |
| Oper Min               |       191100    |       1568.8    |      1   |
| Oper Max               |       196150    |       1528.4    |    102   |
| Oper Value             |       191100    |       1568.8    |      1   |
| Oper Error             |            0    |          0.0    |      0   |
+------------------------+-----------------+-----------------+----------+
| Oper Grid Spacing      |                     50                       |
+------------------------+----------------------------------------------+
```

The following example shows sample output for a port without support for tunability.

```
> port xcvr show port 48 tunability

Not Supported for port 48
```

The following example sets the frequency to 196150 on port 9.

```
> port xcvr set port 9 frequency 196150
```

Procedure 3-17
Setting the port connector mode

Some ports (called dual mode or combination ports) support both RJ45 and SFP (including smaller size SFP+) connectors. Only one of these connectors can be active at a given time. Some ports support a default connector mode, where the mode operates as SFP if a transceiver is installed, or an RJ45 if not. You can set the connector mode manually for the port. Table 3-4 shows the dual mode ports and default mode for each platform that supports them.

Table 3-4
Factory default general port settings by platform

Platform	Ports	Default Connector Mode
3911	9-10	SFP
3930	1-4	Default
3940	1-24	RJ45
5140	1-24	SFP

Note: The port connector mode is applicable only to combination ports that support RJ45 or SFP connectors. It does not apply to XFP ports.

In addition, speed is set to “Auto” with auto-negotiation enabled. So, you can install 1G or 100M transceivers, and the system will automatically set the speed accordingly. If these settings or other port attributes are set explicitly and do not match the capabilities of the active connector, a mismatch warning is generated. The warning is cleared when the attributes match the capabilities of the active connector.

Note: If you attempt to set the mode to a connector that is not supported for the specified port, the system generates a “Capability not supported” error message.

Step	Action

1 Set the mode for a specific port:

```
port set port <port> mode <default|rj45|sfp>
```

where

port <port>	is the port.
mode <default|rj45|sfp>	is the physical interface connector mode.

Copyright© 2012-2013 Ciena® Corporation

Example

The following example sets the mode for port 9 to RJ45.

```
port set port 9 mode rj45
```


Copyright© 2012-2013 Ciena® Corporation

Copyright© 2012-2013 Ciena® Corporation

Hardware resource management

This chapter describes how to configure resource management for customized assignment of hardware resources.



CAUTION
Service disruption
Configuration of resource management requires a reboot to implement changes.

The system assigns hardware resources (classifier, meter, and counter resource types) for various software features. Depending upon the feature, you can reassign these resources to provide additional resources for other features.

This customized resource management is supported for the following features:

- Accelerated CFM over PBB-TE (5150)
- Broadcast containment
- CFM
- DHCP relay
- Traffic profiling
- Virtual circuit statistics
- Virtual switch Layer 2 enhanced transforms
- Transport OAM

Copyright© 2012-2013 Ciena® Corporation

Each resource type is mapped into a number of pools, where each pool contains a number of resources. The number of pools and resources per pool depends upon the platform as shown in Table 4-1.

Table 4-1
Resource pools

Platform	Resource type	Number of pools	Pool boundary
3902, 3911, 3920, 3940, 5140	Classifiers	16	128
	Meters	16	128
	Counters	16	128
3916, 3930, 3931	Classifiers	8	256
	Meters	8	512 (256 meter pairs)
	Counters	8	256
3960, 5150	Classifiers	16	512
	Meters	16	1024 (512 meter pairs)
	Counters	16	512

Meter resources are allocated differently depending upon the platform. On the 3902, 3911, 3920, 3940, and 5140 platforms, the reserved resources limit prevents meters from being consumed by other features. On the 3916, 3930, 3931, 3960 and 5150, the meter resources are global, and availability is not restricted, but is only configurable for certain features. The 3916, 3930, and 3931 platforms have an actual system wide limit of 2048. The 3960 and 5150 platforms have an actual system wide limit of 8192 meters, though the sum of all assigned meter counts can exceed 8192 up to 16,384. The reserved resource value represents a limit to the number of resources for a feature, but does not guarantee availability.

Copyright© 2012-2013 Ciena® Corporation

Each feature consumes hardware resources at a different rate depending upon the configuration as shown in Table 4-2.

Table 4-2
Resources consumed per feature and type

Feature	Resource type	Resources consumed
Accelerated CFM over PBT (configurable on 5150 only)	Classifiers	1
	Meters	0
	Counters	0
Broadcast containment	Classifiers	3 maximum per filter
	Meters	2
	Counters	2
CFM (not configurable on 3916, 3930, 3931)	Classifiers	On 3960 and 5150 platforms, CFM consumes 39 static entries by default. On 3902, 3911, 3920, 3940 and 5140 platforms, CFM consumes classifiers depending upon the configuration and classification type in use. For details, see 009-3203-009, Fault and Performance Management.
	Meters	0
	Counters	0
DHCP relay (not configurable on 3916, 3930, 3931)	Classifiers	On 3960 and 5150 platforms, DHCP relay consumes 2 static entries by default. On 3902, 3911, 3920, 3940 and 5140 platforms, DHCP relay consumes 2 classifiers for each VLAN with DHCP relay enabled.
	Meters	0
	Counters	0

Copyright© 2012-2013 Ciena® Corporation

Table 4-2
Resources consumed per feature and type (continued)

Feature	Resource type	Resources consumed
Traffic profiling	Classifiers	Varies. By default, each port is set to the standard-dot1dpri mode, which consumes 2 classifiers per port. Additional standard traffic profile entries consume classifiers based on the highest number of configured classifiers per type. • standard-dot1dpri - up to 8 • standard-ip-prec - up to 8 • standard-dscp - up to 64 • standard-vlan - 1 • standard-vlan-dot1dpri - up to 9 • standard-vlan-ipp - up to 9 • standard-vlan-dscp - up to 65 • hierarchical-port (3916, 3930, 3931, 3960 and 5150) - Consumes classifiers based upon the sum of the parent and child mode classifiers. • hierarchical-vlan (3916, 3930, 3931, 3960 and 5150) - Consumes classifiers based upon the sum of the parent and child mode classifiers. For example, if you created a traffic profile entry with 2.1D, 2 IP Prec, and 12 DSCP classifiers, the total classifiers consumed would be 12.
	Meters	2
	Counters	2
Transport OAM (configurable on 3916, 3930, 3931, 3960 and 5150 platforms)	Classifiers	250 static entries by default
	Meters	0
	Counters	0

Copyright© 2012-2013 Ciena® Corporation

Table 4-2
Resources consumed per feature and type (continued)

Feature	Resource type	Resources consumed
Virtual circuit statistics	Classifiers	Varies. When statistics are enabled for a virtual circuit, classifiers are consumed depending upon the number of ports in the provider VLAN associated with the virtual circuit and the platform. • On the 3916, 3930, 3931, 3960 and 5150, a virtual circuit consumes 1 classifier for received traffic, and each port in the provider VLAN consumes 1 classifier for transmitted traffic. • On the 3902, 3911, 3920, 3940, 5140, a virtual circuit consumes 2 classifiers for received traffic, and each port in the provider VLAN consumes 2 classifiers for transmitted traffic.
	Meters	0
	Counters	Varies. When statistics are enabled for a virtual circuit, counters are consumed depending upon the number of ports in the provider VLAN associated with the virtual circuit. • On the 3960 and 5150, a virtual circuit consumes 1 counter for received traffic, and each port in the provider VLAN consumes 1 counter for transmitted traffic. • On the 3902, 3911, 3920, 3940, 5140, a virtual circuit consumes 2 counters for received traffic, and each port in the provider VLAN consumes 2 counters for transmitted traffic
Virtual switch enhanced L2 transform (configurable on 3916, 3930, 3931, 3960 and 5150 platforms)	Classifiers	8 static entries by default
	Meters	0
	Counters	0

This chapter provides the following procedures for hardware resources:

- “Configuring resources” on page 4-7
- “Freeing all accelerated CFM over PBB-TE resources” on page 4-9
- “Restoring accelerated CFM over PBB-TE resources to default values” on page 4-11
- “Freeing all broadcast containment resources” on page 4-13
- “Restoring broadcast containment resources to default values” on page 4-14

Copyright© 2012-2013 Ciena® Corporation

- “Freeing CFM resources” on page 4-16
- “Restoring CFM resources to default values” on page 4-17
- “Freeing DHCP relay resources” on page 4-19
- “Restoring DHCP relay resources to default values” on page 4-20
- “Freeing traffic profiling resources” on page 4-22
- “Setting traffic profiling resources” on page 4-24
- “Restoring traffic profiling resources to default values” on page 4-27
- “Freeing virtual circuit statistics resources” on page 4-29
- “Restoring virtual circuit statistics resources to default values” on page 4-30
- “Configuring virtual switch L2 enhanced transform resources” on page 4-32
- “Freeing all virtual switch L2 enhanced transform resources” on page 4-34
- “Configuring transport OAM resources” on page 4-35
- “Freeing all transport OAM resources” on page 4-37
- “Displaying resource configuration information” on page 4-38
- “Resolving resource configuration validation errors” on page 4-39
- “Addressing classifier resource allocation too small for current configuration error” on page 4-40
- “Displaying resource configuration in the configuration file” on page 4-42

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-1
Configuring resources

When setting the pool count, enter the value of the number of resources. The system rounds the number of resources up to the nearest pool boundary listed in Table 4-1 and allocates the specified number of resources by pool. For example, if 750 traffic profile classifier resources are specified to be reallocated on the 3960 platform, 750 is rounded up to the nearest pool boundary, 1024, so 2 pools are reallocated.

On every platform, the number of allocated pools for each resource type must match.

On the 3902, 3911, 3920, 3940, and 5140 platforms, the configured pool count for each resource type must be equal.

On the 3916, 3930, 3931, 3960 and 5150 platforms, for features that use meters, that is, traffic-profiling and broadcast-containment, the configured count for meters must be equal to twice the configured count of the classifier type. For features that do not use meters the value should be 0 for meters. The configured count for the classifier and counter must match, otherwise, the validation will fail.

> ***Note:*** For 3902, 3911, 3920, 3930, and 5140 platforms, the system software requires the reservation of each resource type, regardless of whether the feature uses the resource. For example, traffic profiling and broadcast containment are the only two features that actually consume meter resources. If you are configuring classifiers or counters to reassign them to a feature that doesn't actually consume meters, you still have to configure the assignment of meter resources.

After reassigning resources, you can manually run the command to validate resource configuration. In addition, the system software performs resource validation automatically when you save the configuration. When the validation is successful, the CLI returns to the prompt. If resource validation fails, an error message is generated. For error examples and resolutions, see "Resolving resource configuration validation errors" on page 4-39.

> ***Note:*** If resource validation fails while attempting to save the configuration, the configuration is not saved.

Copyright© 2012-2013 Ciena® Corporation

Once the validation is successful, you need to reboot in order for the resource allocation to take effect.

Step	Action
1	Set the pool counts for each resource type per feature you want to reassign from. `resource-manager pool set resource <classifier|meter|counter> feature <accelerated-cfm-over-pbt | broadcast-containment | cfm | dhcp-relay | traffic-profiling | transport-oam | vc-statistics | vs-enhanced-l2-transform> count <NUMBER>`
2	Set the pool counts for each resource type per feature you want to reassign to. `resource-manager pool set resource <classifier|meter|counter> feature <accelerated-cfm-over-pbt | broadcast-containment | cfm | dhcp-relay | traffic-profiling | transport-oam | vc-statistics | vs-enhanced-l2-transform> count <NUMBER>`
3	Validate the configuration. `resource-manager validate`
4	Save the configuration. `configuration save`
5	Reboot the system to implement the changes. `chassis reboot`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-2
Freeing all accelerated CFM over PBB-TE resources

If accelerated CFM over PBB-TE is not used, you can free up the classifier, meter, and counter resources allocated to accelerated CFM over PBB-TE in order to use them for other features, for example, traffic profiling. Free up classifier, meter, and counter resources by setting them to 0.

CFM over PBB-TE resources are only configurable on the 5150 platform.

Step	Action
1	Disable accelerated CFM over PBB-TE: `cfm service disable service <service>` where service <service> is accelerated CFM over PBB-TE.
2	Unset the accelerate attribute: `cfm service unset service <service> accelerate` where service <service> is accelerated CFM over PBB-TE.
3	Enable accelerated CFM over PBB-TE: `cfm service enable service <service>` where service <service> is accelerated CFM over PBB-TE.
4	Set the pool count for classifier resources to 0. `resource-manager pool set resource classifier feature accelerated-pbt-over-cfm count 0`
5	Set the pool count for meter resources to 0. `resource-manager pool set resource meter feature accelerated-pbt-over-cfm count 0`
6	Set the pool count for counter resources to 0. `resource-manager pool set resource counter feature accelerated-pbt-over-cfm count 0`
7	Validate the configuration. `resource-manager validate`
8	Save the configuration. `configuration save`

Copyright© 2012-2013 Ciena® Corporation

9 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-3
Restoring accelerated CFM over PBB-TE resources to default values

Restore accelerated CFM over PBB-TE resources to default values if you allocated the resources to different features, and then choose to use accelerated CFM over PBB-TE.

Table 4-4 shows the default and maximum number of resources that can be reserved for accelerated CFM over PBB-TE.

Table 4-3
Accelerated CFM over PBB-TE default and maximum resources reservation

Platform	Resource type	Default	Maximum resources
5150	Classifiers	512	512
	Meters	0	512
	Counters	512	512

Step	Action

1 Set the pool count for classifier resources to the default for the platform:

```
resource-manager pool set resource classifier feature
accelerated-pbt-over-cfm count <ClassifierResources>
```

where

<ClassifierResources> is the default pool count for classifier resources for the platform.

2 Set the pool count for meter resources to the default for the platform.

```
resource-manager pool set resource meter feature
accelerated-pbt-over-cfm count <MeterResources>.
```

where

<MeterResources> is the default pool count for meter resources for the platform.

Copyright© 2012-2013 Ciena® Corporation

3 Set the pool count for counter resources to the default for the platform.

```
resource-manager pool set resource counter feature
accelerated-pbt-over-cfm count <CounterResources>.
```

where

<CounterResources> is the default pool count for counter resources for the platform.

4 Validate the configuration.

```
resource-manager validate
```

5 Save the configuration.

```
configuration save
```

6 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-4
Freeing all broadcast containment resources

If broadcast containment is not used, you can free up the classifier, meter, and counter resources allocated to broadcast containment in order to use them for other features, for example, traffic profiling. Free up classifier, meter, and counter resources by setting them to 0.

Resources are not required if the broadcast containment resource mode is set to off. You can set the broadcast containment resource mode by means of the broadcast containment set resource mode off command.

Step	Action
1	Delete any existing broadcast containment filters. `broadcast-containment delete filter <BcastFilterName>` where <BcastFilterName> is the name of the broadcast containment filter to delete
2	Set the pool count for classifier resources to 0. `resource-manager pool set resource classifier feature broadcast-containment count 0`
3	Set the pool count for meter resources to 0. `resource-manager pool set resource meter feature broadcast-containment count 0`
4	Set the pool count for counter resources to 0. `resource-manager pool set resource counter feature broadcast-containment count 0`
5	Validate the configuration. `resource-manager validate`
6	Save the configuration. `configuration save`
7	Reboot the system to implement the changes. `chassis reboot`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-5
Restoring broadcast containment resources to default values

Restore broadcast containment resources to default values if you allocated the resources to different features, and then choose to use broadcast containment.

Table 4-4 shows the default and maximum number of resources that can be reserved for broadcast containment.

Table 4-4
Broadcast containment default and maximum resources reservation per platform

Platform	Resource type	Default	Maximum resources
3902, 3911, 3920, 3940, 5140	Classifiers	128	128
	Meters	128	128
	Counters	128	128
3916, 3930, 3931	Classifiers	256	256
	Meters	512	512
	Counters	256	256
3960 and 5150	Classifiers	512	512
	Meters	1024	1024
	Counters	512	512

Step	Action

1 Set the pool count for classifier resources to the default for the platform:

```
resource-manager pool set resource classifier feature
broadcast-containment count <ClassifierResources>
```

where

<ClassifierResources> is the default pool count for classifier resources for the platform.

Copyright© 2012-2013 Ciena® Corporation

2 Set the pool count for meter resources to the default for the platform.

```
resource-manager pool set resource meter feature
broadcast-containment count <MeterResources>.
```

where

<MeterResources> is the default pool count for meter resources for the platform.

3 Set the pool count for counter resources to the default for the platform.

```
resource-manager pool set resource counter feature
broadcast-containment count <CounterResources>.
```

where

<CounterResources> is the default pool count for counter resources for the platform.

4 Validate the configuration.

```
resource-manager validate
```

5 Save the configuration.

```
configuration save
```

6 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-6
Freeing CFM resources

If for any reason, the number of consumed classifiers exceeds the configured reservation, then CFM will switch from service network mode to global mode. In global mode, classifiers are based on the EtherType and MD level only, rather than service network, EtherType, and MD level, so any frame with a CFM or Y.1731 EtherType will be delivered to the CPU.

If CFM is not used, you can reallocate the resources reserved for classifier, meter, and counter resources to 0 and use them for another feature.

CFM resources are only configurable on the 3902, 3911, 3920, 3940, 3960, 5140 and 5150 platforms.

Step	Action
1	Delete any existing CFM services. `cfm service delete service <CfmService>` where <CfmService> is the CFM service to delete
2	Set the pool count for classifier resources to 0. `resource-manager pool set resource classifier feature cfm count 0`
3	Set the pool count for meter resources to 0. `resource-manager pool set resource meter feature cfm count 0`
4	Set the pool count for counter resources to 0. `resource-manager pool set resource counter feature cfm count 0`
5	Validate the configuration. `resource-manager validate`
6	Save the configuration. `configuration save`
7	Reboot the system to implement the changes. `chassis reboot`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-7
Restoring CFM resources to default values

Restore CFM resources to default values if you allocated the resources to different features, and then choose to use CFM.

Table 4-5 shows the default and maximum number of resources that can be reserved for CFM.

Table 4-5
CFM default and maximum resource reservation per platform

Platform	Resource type	Default	Maximum resources
3902, 3911, 3920, 3940, 5140	Classifiers	128	1792
	Meters	128	1792
	Counters	128	1792
3916, 3930, 3931	Classifiers	Not applicable	Not applicable
	Meters	Not applicable	Not applicable
	Counters	Not applicable	Not applicable
3960 and 5150	Classifiers	512	512
	Meters	0	0
	Counters	512	512

Step	Action

1 Set the pool count for classifier resources to the default for the platform.

```
resource-manager pool set resource classifier feature cfm
count <ClassifierResources>
```

where

<ClassifierResources> is the default pool count for classifier resources for the platform.

Copyright© 2012-2013 Ciena® Corporation

2 Set the pool count for meter resources to the default for the platform.

```
resource-manager pool set resource meter feature cfm
count <MeterResources>
```

where

<MeterResources> is the default pool count for meter resources for the platform.

3 Set the pool count for counter resources to the default for the platform.

```
resource-manager pool set resource counter feature cfm
count <CounterResources>
```

where

<CounterResources> is the default pool count for counter resources for the platform.

4 Validate the configuration.

```
resource-manager validate
```

5 Save the configuration.

```
configuration save
```

6 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-8
Freeing DHCP relay resources

DHCP relay resources are only configurable on the 3902, 3911, 3920, 3940, 3960, 5140 and 5150 platforms.

If DHCP relay is not used, you can reallocate the resources reserved for classifier, meter, and counter resources to 0 and use them for another feature.

Step	Action
1	Delete any existing DHCP relay entries. `dhcp l2-relay-agent delete vlan <VlanList>` where <VlanList> is the VLAN ID of the relay agent
2	Set the pool count for classifier resources to 0. `resource-manager pool set resource classifier feature dhcp-relay count 0`
3	Set the pool count for meter resources to 0. `resource-manager pool set resource meter feature dhcp-relay count 0`
4	Set the pool count for counter resources to 0. `resource-manager pool set resource counter feature dhcp-relay count 0`
5	Validate the configuration. `resource-manager validate`
6	Save the configuration. `configuration save`
7	Reboot the system to implement the changes. `chassis reboot`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-9
Restoring DHCP relay resources to default values

Restore DHCP relay resources to default values if you allocated the resources to different features, and then choose to use DHCP relay.

Table 4-6 shows the default and maximum number of resources that can be reserved for DHCP relay.

Table 4-6
DHCP relay default and maximum resource reservation per platform

Platform	Resource type	Default	Maximum resources
3902, 3911, 3920, 3940, 5140	Classifiers	128	256
	Meters	128	256
	Counters	128	256
3916, 3930, 3931	Classifiers	Not configurable	Not configurable
	Meters	Not configurable	Not configurable
	Counters	Not configurable	Not configurable
3960 and 5150	Classifiers	512	512
	Meters	0	0
	Counters	512	512

Step	Action

1 Set the pool count for classifier resources to the default for the platform.

```
resource-manager pool set resource classifier feature
dhcp-relay count <ClassifierResources>
```

where

<ClassifierResources> is the default pool count for classifier resources for the platform.

Copyright© 2012-2013 Ciena® Corporation

2 Set the pool count for meter resources to the default for the platform.

```
resource-manager pool set resource meter feature dhcp-
relay count <MeterResources>
```

where

<MeterResources> is the default pool count for meter resources for the platform.

3 Set the pool count for counter resources to the default for the platform.

```
resource-manager pool set resource counter feature
dhcp-relay count <CounterResources>
```

where

<CounterResources> is the default pool count for counter resources for the platform.

4 Validate the configuration.

```
resource-manager validate
```

5 Save the configuration.

```
configuration save
```

6 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-10
Freeing traffic profiling resources

If traffic profiling is not used, you can reallocate the resources reserved for classifier, meter, and counter resources to 0 and use them for another feature.

Note: On 3916, 3930, 3931, 3960, and 5150 platforms, which use the global meter pool, the maximum resources shows a configurable maximum and an actual maximum.

— The configurable maximum is twice the number of classifiers and counters and the value can be greater than the global meter pool.

— The actual maximum is equal to the size of the global meter pool.

Before reassigning traffic profiling resources for other features on 3902, 3911, 3920, 3940, and 5140 platforms, you may need to adjust the port to traffic profiling meter pool assignments to make sure there are enough traffic profiling meter pools that are empty. When ports are associated with a traffic profiling meter pool as described in "Setting traffic profiling port attributes" on page 13-43, hardware resources are allocated for classifiers, meters, and counters (statistics). Also, for each associated port, 2 classifier resources, 1 for ARP and 1 for non-conforming standard profiles, are automatically consumed.

Step	Action

1 Delete any existing traffic profiles.

```
traffic-profiling delete port <PortNameList> profile
<TrafficProfStd>
```

where

<PortNameList>	is the list of ports that you want to delete existing traffic profiles from.
<TrafficProfStd>	sets the parent profile.

2 Set the traffic profiling mode to none for each port.

```
traffic-profiling set <PortName> mode none
```

where

<PortName>	is the name of the port.

3 Set the pool count for classifier resources to 0.

```
resource-manager pool set resource classifier feature
traffic-profiling count 0
```


Copyright© 2012-2013 Ciena® Corporation

4 Set the pool count for meter resources to 0.

```
resource-manager pool set resource meter feature traffic-
profiling count 0
```

5 Set the pool count for counter resources to 0.

```
resource-manager pool set resource counter feature
traffic-profiling count 0
```

6 Validate the configuration.

```
resource-manager validate
```

7 Save the configuration.

```
configuration save
```

8 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-11
Setting traffic profiling resources

The default and maximum number of resources that can be reserved for traffic profiling per resource type are shown in Table 4-7.

Table 4-7
Traffic profiling default and maximum resource reservation

Platform	Resource type	Default	Maximum resources
3902	Classifiers	256	256
	Meters	256	256
	Counters	256	256
3911, 3920, 3940, 5140	Classifiers	1024	1792
	Meters	1024	1792
	Counters	1024	1792
3916, 3930, 3931	Classifiers	1024	1536
	Meters	2048	3072 (configurable) 2048 (actual)
	Counters	1024	1536
3960 and 5150	Classifiers	4096	7168
	Meters	8192	14336 (configurable) 8192 (actual)
	Counters	4096	7168

Note: Traffic profiling meter pools are only applicable on 3902, 3911, 3920, 3940, and 5140 platforms.

Copyright© 2012-2013 Ciena® Corporation

The default configuration on the 3920 reserves 8 resource pools for the traffic-profiling feature. Since the default traffic profiling meter pool only has 2 free meter pools, to reassign 3 resource pools, you need to empty 1 traffic profiling meter pool by reassigning the ports in one of the pools to different non-empty traffic profiling meter pools, otherwise the validation will fail.

For 3920, the default number of resource pools is 8, so total resources is 1024 = 128 X 8. To reduce the number of resources by 3 resource pools, set the number of resources to use 5 resource pools, which is 640 = 128 X 5.

Step	Action

1 Determine available non-empty meter pools:

```
traffic-profiling show meter-pool
```

2 Reassign the ports in one of the meter pools to a different non-empty meter pool:

```
traffic-profiling set port <PortName> meter-pool
<MeterPool>
```

where

port <PortName>	is the port to be reassigned.
meter-pool <MeterPool>	is the non-empty meter pool that the ports are to be reassigned to.

3 Set classifier resources to lower the pool count by three resource pools:

```
resource-manager pool set resource classifier feature
traffic-profiling count <NUMBER>
```

where

count <NUMBER>	is the resource count.

4 Set meter resources to lower the pool count by three resource pools:

```
resource-manager pool set resource meter feature
traffic-profiling count <NUMBER>
```

where

count <NUMBER>	is the resource count.

5 Set counter resources to lower the pool count by three resource pools:

```
resource-manager pool set resource counter feature
traffic-profiling count <NUMBER>
```

where

count <NUMBER>	is the resource count.

Copyright© 2012-2013 Ciena® Corporation

6 Validate the configuration:

```
resource-manager validate
```

7 Save the configuration:

```
configuration save
```

8 Reboot the system to implement the changes:

```
chassis reboot
```

—end—

Example

The following example shows sample output for the traffic-profiling show meter-pool command, which is used to determine available non-empty meter pools in step 1.

```
traffic-profiling show meter-pool

+--------------------- TRAFFIC-PROFILING METER-POOL MAP ----------------------+
| Meter-Pool |          Ports          |Class | Used |Meters| Used |Stats | Used |
+------------+-------------------------+------+------+------+------+------+------+
| TP-POOL1   |                         | 128  | 2    | 64   | 0    | 64   | 0    |
+------------+-------------------------+------+------+------+------+------+------+
| TP-POOL2   |                         | 128  | 0    | 64   | 0    | 64   | 0    |
+------------+-------------------------+------+------+------+------+------+------+
| TP-POOL3   | 1 2                     | 128  | 4    | 64   | 0    | 64   | 0    |
+------------+-------------------------+------+------+------+------+------+------+
| TP-POOL4   | 3 4                     | 128  | 4    | 64   | 0    | 64   | 0    |
+------------+-------------------------+------+------+------+------+------+------+
| TP-POOL5   | 5 6                     | 128  | 4    | 64   | 0    | 64   | 0    |
+------------+-------------------------+------+------+------+------+------+------+
| TP-POOL6   | 7 8                     | 128  | 3    | 64   | 1    | 64   | 1    |
+------------+-------------------------+------+------+------+------+------+------+
| TP-POOL7   | 9 10                    | 128  | 4    | 64   | 0    | 64   | 0    |
+------------+-------------------------+------+------+------+------+------+------+
| TP-POOL8   | 11 12                   | 128  | 4    | 64   | 0    | 64   | 0    |
+------------+-------------------------+------+------+------+------+------+------+
```

In this example, ports 3 and 4 from TP-POOL4 are re-assigned to TP-POOL3 and TP-POOL5 respectively.

```
traffic-profiling set port 3 meter-pool TP-POOL3
traffic-profiling set port 4 meter-pool TP-POOL5
```

Procedure 4-12
Restoring traffic profiling resources to default values

Restore traffic profiling resources to default values if you allocated the resources to different features, and then choose to use traffic profiling.

The default and maximum number of resources that can be reserved for traffic profiling per resource type are shown in Table 4-8.

Table 4-8
Traffic profiling default and maximum resource reservation

Platform	Resource type	Default	Maximum resources
3902	Classifiers	256	256
	Meters	256	256
	Counters	256	256
3911, 3920, 3940, 5140	Classifiers	1024	1792
	Meters	1024	1792
	Counters	1024	1792
3916, 3930, 3931	Classifiers	1024	1536
	Meters	2048	3072 (configurable) 2048 (actual)
	Counters	1024	1536
3960 and 5150	Classifiers	4096	7168
	Meters	8192	14336 (configurable) 8192 (actual)
	Counters	4096	7168

Copyright© 2012-2013 Ciena® Corporation

Step	Action

1 Set the pool count for classifier resources to the default for the platform.

```
resource-manager pool set resource classifier feature
traffic-profiling count <ClassifierResources>
```

where

<ClassifierResources> is the default pool count for classifier resources for the platform.

2 Set the pool count for meter resources to the default for the platform.

```
resource-manager pool set resource meter feature
traffic-profiling count <MeterResources>
```

where

<MeterResources> is the default pool count for meter resources for the platform.

3 Set the pool count for counter resources to the default for the platform.

```
resource-manager pool set resource counter feature
traffic-profiling count <CounterResources>
```

where

<CounterResources> is the default pool count for counter resources for the platform.

4 Validate the configuration.

```
resource-manager validate
```

5 Save the configuration.

```
configuration save
```

6 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-13
Freeing virtual circuit statistics resources

If virtual circuit statistics are not used, you can reallocate the resources reserved for classifier, meter, and counter resources to 0 and use them for another feature.

Note: On 3916, 3930, 3931, 3960, and 5150 platforms, which use the global meter pool, the maximum resources shows a configurable maximum and an actual maximum.

— The configurable maximum is twice the number of classifiers and counters and the value can be greater than the global meter pool.

— The actual maximum is equal to the size of the global meter pool.

Step	Action
1	Disable statistics for each virtual circuit with statistics collection turned on. `virtual-circuit ethernet set vc <VcEth> statistics off` where <VcEth> is the virtual circuit.
2	Set the pool count for classifier resources to 0. `resource-manager pool set resource classifier feature vc-statistics count 0`
3	Set the pool count for meter resources to 0. `resource-manager pool set resource meter feature vc-statistics count 0`
4	Set the pool count for counter resources to 0. `resource-manager pool set resource counter feature vc-statistics count 0`
5	Validate the configuration. `resource-manager validate`
6	Save the configuration. `configuration save`
7	Reboot the system to implement the changes. `chassis reboot`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-14
Restoring virtual circuit statistics resources to default values

Restore virtual circuit statistics resources to default values if you allocated the resources to different features, and then choose to use virtual statistics resources.

The default and maximum number of resources that can be reserved for virtual circuit statistics are shown in Table 4-9.

Table 4-9
CFM default and maximum resource reservation per platform

Platform	Resource type	Default	Maximum resources
3902, 3911, 3920, 3940, 5140	Classifiers	128	1792
	Meters	128	1792
	Counters	128	1792
3916, 3930, 3931	Classifiers	128	1536
	Meters	0	0
	Counters	128	1536
3960 and 5150	Classifiers	512	7168
	Meters	0	0
	Counters	512	7168

Step	Action

1 Set the pool count for classifier resources to the default for the platform.

```
resource-manager pool set resource classifier feature
vc-statistics count <ClassifierResources>
```

where

<ClassifierResources> is the default pool count for classifier resources for the platform.

Copyright© 2012-2013 Ciena® Corporation

2 Set the pool count for meter resources to the default for the platform.

```
resource-manager pool set resource meter feature
vc-statistics count <MeterResources>
```

where

<MeterResources> is the default pool count for meter resources for the platform.

3 Set the pool count for counter resources to the default for the platform.

```
resource-manager pool set resource counter feature
vc-statistics count <CounterResources>
```

where

<CounterResources> is the default pool count for meter resources for the platform.

4 Validate the configuration.

```
resource-manager validate
```

5 Save the configuration.

```
configuration save
```

6 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-15
Configuring virtual switch L2 enhanced transform resources

Virtual switch L2 enhanced transform resources are only applicable to the 3916, 3930, 3931, 3960 and 5150 platform. Attempting to configure virtual switch L2 enhanced transform resources on the 3902, 3911, 3920, 3940, and 5140 platforms will generate an error.

By default, virtual switch L2 enhanced transform does not have any assigned resources. To use this feature, you need to assign available classifier and counter resources to it.

The default and maximum number of resources that can be reserved for virtual switch L2 enhanced transform are shown in Table 4-10.

Table 4-10
Virtual switch L2 enhanced transform default and maximum resource reservations per platform

Platform	Resource type	Default	Maximum resources
3902, 3911, 3920, 3940, 5140	Classifiers	Not applicable	Not applicable
	Meters	Not applicable	Not applicable
	Counters	Not applicable	Not applicable
3916, 3930, 3931	Classifiers	0	256
	Meters	0	0
	Counters	0	256
3960, 5150	Classifiers	0	512
	Meters	0	0
	Counters	0	512

Step	Action
1	Set the pool count for classifier resources. `resource-manager pool set resource classifier feature vs-enhanced-l2-transform count <NUMBER>` where count <NUMBER> is the number of resources to assign.

Copyright© 2012-2013 Ciena® Corporation

2 Set the pool count for counter resources.

```
resource-manager pool set resource counter feature
vs-enhanced-l2-transform count <CounterResources>
```

where

<CounterResources> is the pool count for counter resources for the platform.

3 Validate the configuration.

```
resource-manager validate
```

4 Save the configuration.

```
configuration save
```

5 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-16
Freeing all virtual switch L2 enhanced transform resources

If virtual switch L2 enhanced transform is no longer used, you can reallocate the resources reserved for classifier and counter resources back to 0 and use them for another feature.

Note that 0 is the default value for classifier and counter resources.

Step	Action
1	Remove the ports with virtual switch L2 enhanced transform configuration from their associated virtual switch. `virtual-switch ethernet delete port <PortNameList> vlan <VlanList>` where <VlanList> is the virtual switch.
2	Set the virtual switch L2 transform mode to the default for all ports. `port set port 1-12 vs-l2-transform i-push,e-pop`
3	Set the pool count for classifier resources to 0. `resource-manager pool set resource classifier feature vs-enhanced-l2-transform count 0`
4	Set the pool count for counter resources to 0. `resource-manager pool set resource counter feature vs-enhanced-l2-transform count 0`
5	Validate the configuration. `resource-manager validate`
6	Save the configuration. `configuration save`
7	Reboot the system to implement the changes. `chassis reboot`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-17
Configuring transport OAM resources

Configure transport OAM resources as required for use with MPLS features such as CFM over MPLS, VCCV, LSP BFD, and VS-based remote management.

Transport OAM resources are only applicable to the 3916, 3930, 3931, 3960 and 5150 platforms.

On the 3916, 3930, and 3931 platforms, transport OAM does not have any assigned resources by default. To use this feature, you need to assign available classifier and counter resources to it.

The default and maximum number of resources that can be reserved for transport OAM are shown in Table 4-11.

Table 4-11
Transport OAM default and maximum resource reservations per platform

Platform	Resource type	Default	Maximum resources
3902, 3911, 3920, 3940, 5140	Classifiers	Not applicable	Not applicable
	Meters	Not applicable	Not applicable
	Counters	Not applicable	Not applicable
3916, 3930, 3931	Classifiers	0	256
	Meters	0	0
	Counters	0	256
3960, 5150	Classifiers	512	512
	Meters	0	0
	Counters	512	512

Step	Action
1	Set the pool count for classifier resources. `resource-manager pool set resource classifier feature transport-oam count <NUMBER>` where count <NUMBER> is the number of resources to assign.

Copyright© 2012-2013 Ciena® Corporation

2 Set the pool count for counter resources.

```
resource-manager pool set resource counter feature
transport-oam count <CounterResources>
```

where

<CounterResources> is the pool count for counter resources for the platform.

3 Validate the configuration.

```
resource-manager validate
```

4 Save the configuration.

```
configuration save
```

5 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Procedure 4-18
Freeing all transport OAM resources

If transport OAM is no longer used, you can reallocate the resources reserved for classifier and counter resources back to 0 and use them for another feature.

Note that 0 is the default value for classifier and counter resources on the 3916, 3930, and 3931 platforms. The default value for classifier, meter, and counter resources on the 3960 and 5150 platforms is 512.

Step	Action
1	Set the pool count for classifier resources to 0. `resource-manager pool set resource classifier feature transport-oam count 0`
2	Set the pool count for counter resources to 0. `resource-manager pool set resource counter feature transport-oam count 0`
3	Validate the configuration. `resource-manager validate`
4	Save the configuration. `configuration save`
5	Reboot the system to implement the changes. `chassis reboot`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-19
Displaying resource configuration information

The system software provides a way to display the *active* configuration (what is currently running) and the *candidate* configuration (what it will be upon successful configuration save and reboot).You can also display detailed resource information for the CFM feature on 3902, 3911, 3920, 3940, and 5140 platforms and for the traffic profiling feature on all platforms.

Step	Action

1 Determine the resource configuration information that you want to display.

If you want to display	Then
active and candidate information	Perform step 2.
active configuration only	Perform step 3.
candidate configuration only	Perform step 4.
detailed resource information for the CFM feature on the 3902, 3911, 3920, 3940, and 5140 platforms	Perform step 5.
detailed resource information for the traffic profiling feature	Perform step 6.

2 Display both active and candidate configuration:

```
resource-manager show
```

3 Display active configuration only:

```
resource-manager show active
```

4 Display candidate configuration only:

```
resource-manager show candidate
```

5 Display detailed resource information for the CFM feature on the 3902, 3911, 3920, 3940, and 5140 platforms:

```
resource-manager show feature cfm
```

6 Display detailed resource information for the traffic profiling feature:

```
resource-manager show feature traffic-profiling
```

—end—

Procedure 4-20
Resolving resource configuration validation errors

This section provides examples of some common validation error messages and the method to resolve them.

Example - ERROR: <feature>: classifier, meter, and counter pool counts must match:

```
resource-manager validate

ERROR: traffic-profiling: classifier, meter, and counter pool counts must match
ERROR:                  : classifier pools=14 meter pools=13 counter pools=14
ERROR: Resource validation failed: Resource allocation invalid
```

Step Action

1 Display the candidate configuration.

```
resource-manager show candidate

+--------------------------+------------+------------+------------+------------+----------+
|                          CANDIDATE RESOURCE CONFIGURATION                               |
+--------------------------+------------+------------+------------+------------+----------+
| feature                  |   resource |   reserved |       used |  increment |pltfm max |
+--------------------------+------------+------------+------------+------------+----------+
...
```

2 Scroll to the section for the feature listed in the error message to find the candidate configuration for the resource type.

```
...
| traffic-profiling        | classifier |       1792 |         20 |        128 |       1792 |
|                          |      meter |        512 |          0 |        128 |       1792 |
|                          |    counter |       1792 |          0 |        128 |       1792 |
...
```

3 Set the pool count for the resource types to match.

Note: On the 3916, 3930, 3931, 3960 and 5150 platforms, in order for the pool counts to match, the reserved meter resources must be twice that of the reserved classifier resources. Also, the classifier and counter reserved resources must be equal.

```
resource-manager pool set resource meter feature
traffic-profiling count 1792
```

4 Re-validate the configuration.

```
resource-manager validate
```

5 Save the configuration.

```
configuration save
```

6 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 4-21
Addressing classifier resource allocation too small for current configuration error

If there are not enough resources to support the configuration of a feature, the following error is displayed when the resource-manager validate command is executed:

```
ERROR: Resource validation failed: Resource allocation invalid
```

This error condition is addresed by performing one of the following:

- increasing the pool count so that the number of resources matches the active configuration
- modifying the configuration of the feature

Step	Action

To increase the pool count so that the number of resources matches the active configuration

1 Set the pool count for classifier, meter, and counter resources to the match the active value shown in the error message.

```
resource-manager pool set resource classifier feature
vc-statistics count <ActiveValue>
```

```
resource-manager pool set resource meter feature
vc-statistics count <ActiveValue>
```

```
resource-manager pool set resource counter feature
vc-statistics count <ActiveValue>
```

where

<ActiveValue> is the active value displayed in the error message

2 Re-validate the configuration.

```
resource-manager validate
```

3 Save the configuration.

```
configuration save
```

4 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Step	Action

To modify the configuration of the feature

Copyright© 2012-2013 Ciena® Corporation

1 Display the configuration of the feature. In this case, check for virtual circuits with statistics collection turned on.

```
virtual-circuit show
+------ ETHERNET VIRTUAL CIRCUIT TABLE -----+
| Name             | VLAN          |  Stats   |
+------------------+---------------------------+
| 101              | 101           | On       |
+------------------+---------------------------+
```

2 Modify the configuration of the feature.

```
virtual-circuit ethernet set vc 101 statistics off
```

3 Re-validate the configuration.

```
resource-manager validate
```

4 Save the configuration.

```
configuration save
```

5 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Example

The following example shows sample output for the resource-manager validate command for a configuration where there are not enough resources to support the configuration of the virtual circuit statistics feature.

```
resource-manager validate
ERROR: vc-statistics: classifier resource allocation too
small for current configuration
ERROR:                : candidate=0 active=128 used=2
ERROR: vc-statistics: counter resource allocation too small
for current configuration
ERROR:                : candidate=0 active=128 used=2
ERROR: Resource validation failed: Resource allocation
invalid
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 4-22
Displaying resource configuration in the configuration file

Once resource validation passes and the configuration file is saved, resource configuration is saved in the RESOURCE CONFIG section in the configuration file.

Step	Action
1	Vew resource configuration in the configuration file: `configuration show`

—end—

Example

The following example shows sample output from the configuration show command.

```
configuration show

! CN 3911 Configuration File
! Chassis MAC:   00:02:a1:22:bd:40
! Created:       Thu Oct 22 21:38:31 2009
! Created by:    CLI
! SW Package:    Slot 1 - saos-06-06-00-0103
! Build Number:  4708
! MIB Number:    02-03-11-0027
!
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
! RESOURCE CONFIG
!
resource-manager pool set feature traffic-profiling resource classifier count 1792
resource-manager pool set feature traffic-profiling resource meter count 1792
resource-manager pool set feature traffic-profiling resource counter count 1792
resource-manager pool set feature broadcast-containment resource classifier count 0
resource-manager pool set feature broadcast-containment resource meter count 0
resource-manager pool set feature broadcast-containment resource counter count 0
resource-manager pool set feature cfm resource classifier count 0
resource-manager pool set feature cfm resource meter count 0
resource-manager pool set feature cfm resource counter count 0
resource-manager pool set feature vc-statistics resource classifier count 0
resource-manager pool set feature vc-statistics resource meter count 0
resource-manager pool set feature vc-statistics resource counter count 0
resource-manager pool set feature dhcp-relay resource classifier count 0
resource-manager pool set feature dhcp-relay resource meter count 0
resource-manager pool set feature dhcp-relay resource counter count 0
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
[more 9%] (q,g,space,enter)
```


Copyright© 2012-2013 Ciena® Corporation

Synchronization configuration

This chapter describes system timing and synchronization configuration.

System timing is provided by means of:

- "Synchronous Ethernet"
- "Building Integrated Timing Supply"

This chapter also provides "Network configuration examples".

Synchronous Ethernet

Synchronous Ethernet (SyncE) support provides a migration path from existing frequency synchronization distribution architectures based on SONET/SDH or GPS to a next generation packet network-based frequency synchronization architecture based on Carrier Ethernet with SyncE.

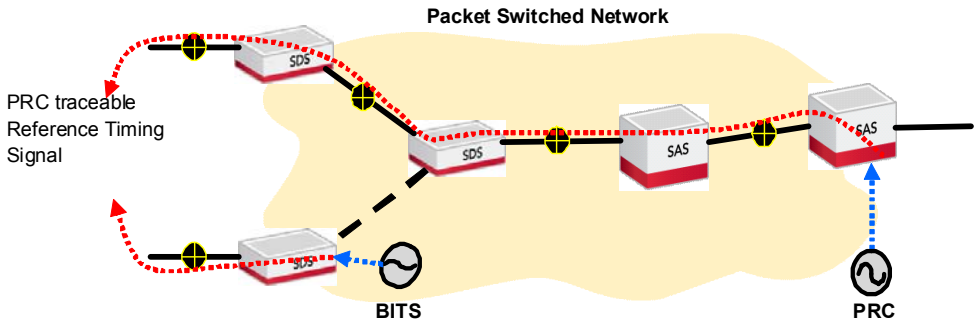
SyncE assures that frequency is distributed at the physical layer where it is not subject to load impairments such as packet congestion and loss.

There can be a number of Ethernet switches involved in the distribution of the reference timing signal. In such cases, the synchronization function within these Ethernet switches must be able to recover synchronization line timing from the incoming bit stream and propagate synchronization line timing to outgoing Ethernet ports.

Figure 5-1 shows a reference timing signal, traceable to a Primary Reference Clock (PRC), which is injected into the Ethernet switch using an external clock port. This signal is extracted and processed by means of a synchronization function before injecting timing onto the Ethernet bit stream leaving the Ethernet switch. The clock supporting synchronous Ethernet networks is called the Ethernet Equipment Clock (EEC).

Copyright© 2012-2013 Ciena® Corporation

Figure 5-1
Example of a primary-secondary synchronization network over synchronous Ethernet



Synchronization Status Messaging (SSM) is used to convey priority and clock traceability. SSM messages indicate the quality level of the system clocks located in the various network elements. Quality level is the holdover performance of a clock.

Clock selection algorithm

The clock selection algorithm applies for input references added to a frequency protection-group. The protection-group selects and locks to the “best” input reference based on the following criteria (in decreasing order of significance):

- forced-reference
- reference failure from
  - — loss of signal
  - — loss of frame
  - — alarm indication signal
  - — hardware faults
  - — quality level below threshold
  - — quality level receive timeout
- override priority of reference
- operational quaity level of reference
- preferred reference: a tie scenario, that is, operational quality levels are the same and no priorities are configured
  - — configured preferred-reference is selected
  - — by default BITS reference is preferred over SyncE reference
  - — tie case between SyncE references only: lower port is preferred

Copyright© 2012-2013 Ciena® Corporation

Building Integrated Timing Supply

Building Integrated Timing Supply (BITS) is a timing signal that is used to distribute frequency synchronization in a telecom environment, typically in a central office. It is usually carried over T1/E1 lines. Timing is encoded within the transmitted data signal to synchronize the whole network.

The 3930 Sync + External Timing platform and 5150 Standard with BITS platform input and output standard electrical specifications for a BITS frequency as follows:

- 2.048Mbps E1 compliant to ITU-T G.703-9
- 1.544Mbps T1 compliant to GR-499 and ITU-T G.703-5

Note: The BITS interface can be configured to be either BITS-IN or BITS-OUT. It cannot do both BITS-IN and BITS-OUT simultaneously.

Network configuration examples

Figure 5-2 shows a linear topology example for system timing. In this example, input references are configured on the ports with arrows. The input references lock to the BITS signal and distribute the timing to the SyncE portion of the network.

Figure 5-2
Linear topology

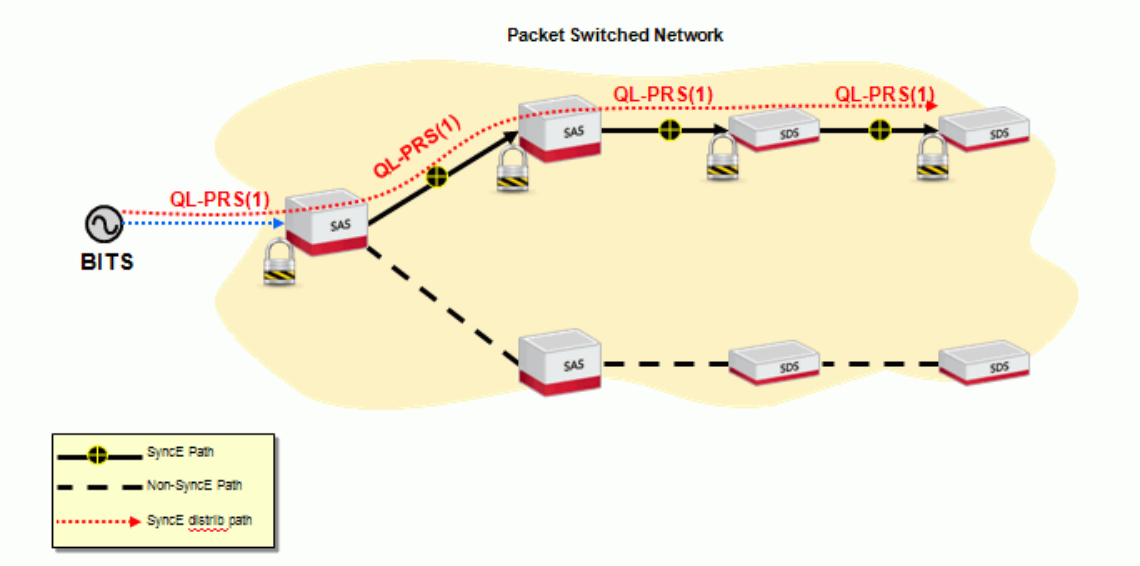
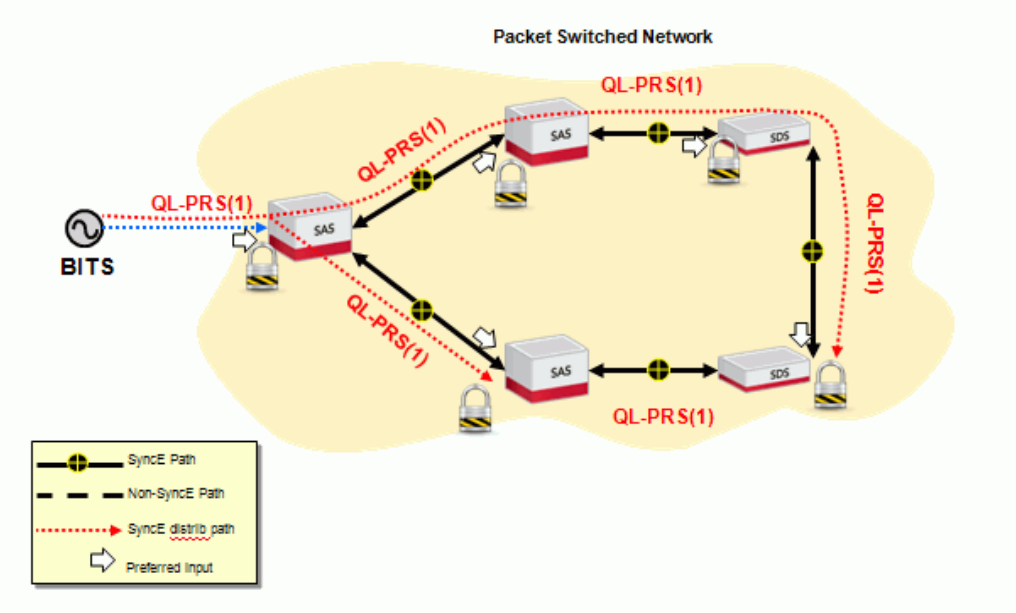


Figure 5-3 shows a ring topology example for system timing. In this example, ports can be configured for input or output. BITS feeds timing to the network. Note that the timing signal does not flow all the way around the ring. SyncE

alone cannot detect or prevent timing loops. Timing loops can only be prevented with careful network planning, that is, proper configuration of preferred inputs and the way the nodes are interconnected. Having a break in the flow avoids timing loops.

Figure 5-3
Ring topology



This chapter provides the following procedures for synchronization:

- “Enabling and disabling synchronization” on page 5-5
- “Configuring synchronization” on page 5-6
- “Configuring frequency input references” on page 5-8
- “Configuring frequency output references” on page 5-12
- “Configuring frequency protection groups” on page 5-15
- “Creating an output port for SyncE from a BITS input” on page 5-18
- “Creating a BITS output from a SyncE input” on page 5-20
- “Displaying information for synchronization” on page 5-22

Copyright© 2012-2013 Ciena® Corporation

Procedure 5-1
Enabling and disabling synchronization

You can

- enable synchronization
- disable synchronization

Step	Action

To enable synchronization

1 Enable synchronization:

```
sync enable
```

To disable synchronization

2 Disable synchronization:

```
sync disable
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 5-2
Configuring synchronization

You can configure

- option type
- reversion mode
- global wait-to-restore timer

Step	Action

To set the option type

1 Set the option type:

```
sync set option-type <option1|option2|none>
```

where

option-type <option1| option2|none>	is the option type. The default value is none. • option1 applies to SDH networks optimized for the 2048 kbit/s hierarchy • option2 applies to SDH networks optimized for the 1544 kbit/s hierarchy that includes the rates 1544 kbit/s, 6312 kbit/s and 44736 kbit/s ***Note:*** When option-type is set to option-1, reversion-mode defaults to revertive. When option-type is set to option-2, reversion-mode defaults to non-revertive. You can explicitly configure reversion-mode. To change the option-type, ensure there are no existing configurations.

To set the reversion-mode

2 Set the global reversion-mode:

```
sync set reversion-mode <non-revertive|revertive>
```

where

reversion-mode <non-revertive| revertive>	is the global reversion mode.

Copyright© 2012-2013 Ciena® Corporation

To set the global wait-to-restore timer

3 Set the global wait-to-restore timer:

```
sync set wait-to-restore <MINUTES: 0-12>
```

where

wait-to-restore <MINUTES: 0-12>	is the time assigned to the wait-to-restore timer. The wait-to-restore timer is used when the reversion mode is set to revertive.

To restore the option type to the default value

4 Restore the option type to the default value:

```
sync unset option-type
```

To restore the reversion mode to the default value

5 Restore the reversion mode to the default value:

```
sync unset reversion-mode
```

To restore the global wait-to-restore timer to the default value

6 Restore the global wait-to-restore timer to the default value:

```
sync unset wait-to-restore
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 5-3
Configuring frequency input references

You can

- create a BITS frequency input reference
- create a SyncE frequency input reference
- delete a frequency input reference
- set attributes of a BITS frequency input reference
- set attributes of a SyncE frequency input reference
- restore attributes of a BITS frequency input reference to default values
- restore attributes of a SyncE frequency input reference to default values

Step	Action

To create a BITS frequency input reference

1 Create a BITS frequency input reference:

```
sync input freq-ref create bits-freq-ref <STRING[32]>
{bits-interface <bits-interface-name>} {mode <e1|t1>}
[encoding ami|b8zs|hdb3] [forced-ql <INTEGER: 0-15> ] [
format e1-crc|e1-no-crc|esf|sf] [override-priority
<INTEGER: 1-10>] [ql-receive disable|enable] [ssm-
location sa4|sa5|sa6|sa7]
```

where

bits-freq-ref <STRING[32]>	is the identifier for the BITS frequency input reference.
bits-interface <bits-interface-name>	is the BITS interface name. Valid values are: • bits-a • bits-b For the 3930 Sync + External Timing platform, only bits-a is used. For the 5150 Standard with BITS platform, assign the BITS interface name that corresponds to the location of the BITS module that is to be activated.
mode <e1| t1>	is the mode.
encoding ami|b8zs|hdb3	is the encoding type.
forced-ql <INTEGER: 0-15>	sets the quality level at which the received quality level is overridden.

Copyright© 2012-2013 Ciena® Corporation

where

format e1-crc|e1-no-crc|esf|sf	is the format.
override-priority <INTEGER: 1-10>	is the override-priority for the BITS frequency input reference.
ql-receive disable|enable	enables or disables received quality level.
ssm-location sa4|sa5|sa6|sa7	is the location for SSM messaging. Ssm-location is only applicable for e1 mode.

To create a SyncE frequency input reference

2 Create a SyncE frequency input reference:

```
sync input freq-ref create synce-freq-ref <STRING[32]>
{port <phy-port-name>} [forced-ql <INTEGER: 0-15>]
[override-priority <INTEGER: 1-10>] [ql-receive
disable|enable]
```

where

synce-freq-ref <STRING[32]>	is the identifier for the SyncE frequency input reference.
port <phy-port-name>	is the physical port.
forced-ql <INTEGER: 0-15>	sets the quality level at which the received quality level is overridden.
override-priority <INTEGER: 1-10>	is the override-priority for the SyncE frequency input reference.
ql-receive disable|enable	enables or disables received quality level.

To delete a frequency input reference

3 Delete a frequency input reference:

```
sync input freq-ref delete freq-ref <STRING[32]>
```

where

freq-ref <STRING[32]>t	is the identifier for the frequency input reference that you want to delete.

Copyright© 2012-2013 Ciena® Corporation

To set attributes of a BITS frequency input reference

4 Set attributes of a BITS frequency input reference:

```
sync input freq-ref set bits-freq-ref <STRING[32]>
{[encoding ami|b8zs|hdb3] [forced-ql <INTEGER: 0-15>]
[override-priority <INTEGER: 1-10>] [ql-receive
disable|enable]} [ssm-location sa4|sa5|sa6|sa7]
```

where

bits-freq-ref <STRING[32]>	is the identifier for the BITS frequency input reference.
encoding ami|b8zs|hdb3	is the encoding type.
forced-ql <INTEGER: 0-15>	sets the quality level at which the received quality level is overridden.
override-priority <INTEGER: 1-10>	is the override-priority for the BITS frequency input reference.
ql-receive disable|enable]	enables or disables received quality level.
ssm-location sa4|sa5|sa6|sa7	is the location for SSM messaging. Ssm-location is only applicable for e1 mode.

To set attributes of a SyncE frequency input reference

5 Set attributes of a SyncE frequency input reference:

```
sync input freq-ref set synce-freq-ref <STRING[32]>
{[forced-ql <INTEGER: 0-15>] [override-priority <INTEGER:
1-10>] [ql-receive disable|enable]}
```

where

synce-freq-ref <STRING[32]>	is the identifier for the SyncE frequency input reference.
forced-ql <INTEGER: 0-15>	sets the quality level at which the received quality level is overridden.
override-priority <INTEGER: 1-10>	is the override-priority for the SyncE frequency input reference.
ql-receive disable|enable]	enables or disables received quality level.

Copyright© 2012-2013 Ciena® Corporation

To restore attributes of a BITS frequency input reference to default values

6 Restore attributes of a BITS frequency input reference to default values:

```
sync input freq-ref unset bits-freq-ref <STRING[32]>
{[encoding] [forced-ql] [override-priority] [ql-receive]
[ssm-location]}
```

where

bits-freq-ref <STRING[32]>	is the identifier for the BITS frequency input reference.
encoding	is the encoding type. Default value is AMI for T1 mode and HDB3 for E1 mode.
forced-ql	is the quality level at which the received quality level is overridden.
override-priority	is the override-priority for the BITS frequency input reference.
ql-receive	enables or disables received quality level. Default value is enabled for E1, that is, e1-crc format, and T1, that is esf format. For all other formats, the default is disabled.
ssm-location	is the location for SSM messaging. This attribute only applies for E1 mode. Default value is SA4.

To restore attributes of a SyncE frequency input reference to default values

7 Restore attributes of a SyncE frequency input reference to default values:

```
sync input freq-ref unset synce-freq-ref <STRING[32]>
{[forced-ql] [override-priority] [ql-receive]}
```

where

synce-freq-ref <STRING[32]>	is the identifier for the SyncE frequency input reference.
forced-ql	is the quality level at which the received quality level is overridden.
override-priority	is the override-priority for the SyncE frequency input reference.
ql-receive	enables or disables received quality level. Default value is enable.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 5-4
Configuring frequency output references

You can

- create a BITS frequency output reference
- create a SyncE frequency output reference
- delete a frequency output reference
- set attributes of a BITS frequency output reference
- restore attributes of a BITS frequency output reference to default values

Step	Action

To create a BITS frequency output reference

1 Create a BITS frequency output reference:

```
sync output freq-ref create bits-freq-ref <STRING[32]>
{bits-interface <bits-interface-name>} {mode e1|t1}
{format e1-crc|e1-no-crc|esf|sf} [encoding
ami|b8zs|hdb3] [line-build-out 133|266|399|533|655]
[ssm-location sa4|sa5|sa6 |sa7]
```

where

bits-freq-ref <STRING[32]>	is the BITS frequency output reference.
bits-interface <bits-interface-name>	is the BITS interface.
mode e1|t1	is the mode. The default value for the encoding attrribute when T1 mode is selected is ami. The default value for associated attributes when E1 mode is selected are: • encoding hdb3 • ssm-location sa4 • line-build-out 133
format e1-crc| e1-no-crc|esf|sf	is the format.

Copyright© 2012-2013 Ciena® Corporation

where

encoding ami|b8zs|hdb3	is the encoding type.
line-build-out 133|266|399|533|655	is the line build-out. For line build-out of 0 to 133 feet, select 133. For line build-out of 133 to 266 feet, select 266. For line-build-out of 266 to 399 feet, select 399. For line-build-out of 399 to 533 feet, select 533. For line-build-out of 533 to 655 feet, select 655.
ssm-location sa4|sa5|sa6 |sa7	is the location for SSM messaging. SSM-location is only applicable for e1 mode.

To create a SyncE frequency output reference

2 Create a SyncE frequency output reference:

```
sync output freq-ref create synce-freq-ref <STRING[32]>
{port <port-id>}
```

where

synce-freq-ref <STRING[32]>	is the SyncE frequency output reference.
port <port-id>	is the port.

To delete a frequency output reference

3 Delete a frequency output reference:

```
sync output freq-ref delete freq-ref <STRING[32]>
```

where

freq-ref <STRING[32]>	is the frequency output reference that you want to delete.

Copyright© 2012-2013 Ciena® Corporation

To set attributes of a BITS frequency output reference.

4 Set attributes of a BITS frequency output reference:

```
sync output freq-ref set bits-freq-ref <STRING[32]>
{[encoding ami|b8zs|hdb3] [line-build-out
133|266|399|533|655] [ssm-location sa4|sa5|sa6|sa7]}
```

where

bits-freq-ref <STRING[32]>	is the BITS frequency output reference.
encoding ami|b8zs|hdb3	is the encoding type.
line-build-out 133|266|399|533|655	is the line build-out. For line build-out of 0 to 133 feet, select 133. For line build-out of 133 to 266 feet, select 266. For line-build-out of 266 to 399 feet, select 399. For line-build-out of 399 to 533 feet, select 533. For line-build-out of 533 to 655 feet, select 655.
ssm-location sa4|sa5|sa6|sa7]	is the location for SSM messaging.

To restore attributes of a BITS frequency output reference to default values

5 Restore attributes of a BITS frequency output reference to default values:

```
sync output freq-ref unset bits-freq-ref <STRING[32]>
{[encoding] [line-build-out] [ssm-location]}
```

where

bits-freq-ref <STRING[32]>	is the BITS frequency output reference.
encoding	is the encoding type.
line-build-out	is the line build-out.
ssm-location	is the location for SSM messaging.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 5-5
Configuring frequency protection groups

You can

- create a frequency protection group
- delete a frequency protection group
- enable a frequency protection group
- disable a frequency protection group
- add a frequency input object to a frequency protection group
- remove a frequency input object from a frequency protection group
- set attributes of a frequency protection group
- restore attributes of a frequency protection group to default values

Step	Action

To create a frequency protection group

1 Create a frequency protection group:

```
sync protection-group create freq-group <STRING[32]>
{threshold-ql <INTEGER: 0-14>}
```

where

freq-group <STRING[32]>	is the frequency protection group.
threshold-ql <INTEGER: 0-14>	is the quality level threshold.

To delete a frequency protection group

2 Delete a frequency protection group:

```
sync protection-group delete freq-group <STRING[32]>
```

where

freq-group <STRING[32]>	is the frequency protection group.

Copyright© 2012-2013 Ciena® Corporation

To enable a frequency protection group

3 Enable a frequency protection group:

```
sync protection-group enable freq-group <STRING[32]>
```

where

freq-group <STRING[32]>	is the frequency protection group.

To disable a frequency protection group

4 Disable a frequency protection group:

```
sync protection-group disable freq-group <STRING[32]>
```

where

freq-group <STRING[32]>	is the frequency protection group.

To add a frequency input object to a frequency protection group

5 Add a frequency input object to a frequency protection group:

```
sync protection-group add freq-group <STRING[32]> {freq-
ref <STRING[32]>}
```

where

freq-group <STRING[32]>	is the frequency protection group.
freq-ref <STRING[32]>	is the frequency input object.

To remove a frequency input object from a frequency protection group

6 Remove a frequency input object from a frequency protection group:

```
sync protection-group remove <freq-group <STRING[32]>
{freq-ref <STRING[32]>}
```

where

freq-group <STRING[32]>	is the frequency protection group.
freq-ref <STRING[32]>	is the frequency input object.

Copyright© 2012-2013 Ciena® Corporation

To set attributes of a frequency protection group

7 Set attributes of a frequency protection group:

```
sync protection-group set freq-group <STRING[32]>
{[forced-ref <STRING[32]>][preferred-ref
<STRING[32]>][threshold-ql <INTEGER: 0-14>]}
```

where

freq-group <STRING[32]>	is the frequency protection group.
forced-ref <STRING[32]>	is the forced reference.
preferred-ref <STRING[32]>	is the preferred reference.
threshold-ql <INTEGER: 0-14>	is the quality level threshold.

To restore attributes of a frequency protection group to default values

8 Restore attributes of a frequency protection group to default values:

```
sync protection-group unset freq-group <STRING[32]>
{[forced-ref][preferred-ref]}
```

where

freq-group <STRING[32]>	is the frequency protection group.
forced-ref	is the forced reference.
preferred-ref	is the preferred reference.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 5-6
Creating an output port for SyncE from a BITS input

Create an output port for SyncE from a BITS input when you want to establish SyncE service based on the reference entering the device at the BITS port. The reference is passed to one or more ports on the device that are connected to other devices.

The BITS connection can be used for a timing reference input or output, but not both at the same time.

Additional outputs can be created as needed on other ports. Each output must have a unique name.

Step	Action
1	Create the BITS input. `sync input freq-ref create bits-freq-ref <STRING[32]> bits-interface <bits-interface-name> mode <e1|t1> format <e1-crc|e1-no-crc|esf|sf>` where bits-freq-ref <STRING[32]> — is the identifier for the BITS frequency input reference. bits-interface <bits-interface-name> — is the BITS interface name. Valid values are: • bits-a • bits-b For the 3930 Sync + External Timing platform, only bits-a is used. For the 5150 Standard with BITS platform, assign the BITS interface name that corresponds to the location of the BITS module that is to be activated. mode <e1|t1> — is the mode. format <e1-crc|e1-no-crc|esf|sf> — is the format.

Copyright© 2012-2013 Ciena® Corporation

2 Create a protection group with a defined threshold quality level (QL).

```
sync protection-group create freq-group <STRING[32]>
threshold-ql <INTEGER: 0-14>
```

where

freq-group <STRING[32]>	is the frequency protection group.
threshold-ql <INTEGER: 0-14>	is the quality level threshold.

3 Add the BITS input frequency reference to the protection group.

```
sync protection-group add freq-group <STRING[32]> freq-
ref <STRING[32]>
```

where

freq-group <STRING[32]>	is the frequency protection group.
freq-ref <STRING[32]>	is the BITS input frequency reference.

4 Create an output port for SyncE.

```
sync output freq-ref create synce-freq-ref <STRING[32]>
{port <port-id>}
```

where

synce-freq-ref <STRING[32]>	is the SyncE frequency output reference.
port <port-id>	is the port.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 5-7
Creating a BITS output from a SyncE input

Create a BITS output from a SyncE input according to system timing requirements.

Step	Action

1 Create the SyncE input frequency reference:

```
sync input freq-ref create synce-freq-ref <STRING[32]>
forced-ql <INTEGER: 0-15> port <phy-port-name>
```

where

synce-freq-ref <STRING[32]>	is the identifier for the SyncE frequency input reference.
port <phy-port-name>	is the physical port.
forced-ql <INTEGER: 0-15>	sets the quality level at which the received quality level is overridden.

2 Create the BITS output frequency reference.

```
sync output freq-ref create bits-freq-ref <STRING[32>
bits-interface <bits-interface-name> mode <e1|t1> format
<e1-crc|e1-no-crc|esf|sf>
```

where

bits-freq-ref <STRING[32]>	is the BITS frequency output reference.
bits-interface <bits-interface-name>	is the BITS interface.
mode <e1|t1>	is the mode. The default value for the encoding attrribute when T1 mode is selected is ami. The default value for associated attributes when E1 mode is selected are: • encoding hdb3 • ssm-location sa4 • line-build-out 133
format <e1-crc|e1-no-crc|esf|sf>	is the format.

Copyright© 2012-2013 Ciena® Corporation

3 Create the protection group:

```
sync protection-group create freq-group <STRING[32]>
threshold-ql <INTEGER: 0-14>
```

where

freq-group <STRING[32]>	is the frequency protection group.
threshold-ql <INTEGER: 0-14>	is the quality level threshold.

4 Add the SyncE input frequency reference to the protection group:

```
sync protection-group add freq-group <STRING[32]> freq-
ref <STRING[32]>
```

where

freq-group <STRING[32]>	is the frequency protection group.
freq-ref <STRING[32]>	is the frequency reference.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 5-8
Displaying information for synchronization

You can display

- global synchronization information
- quality level hierarchy table, as defined in G.781
- overall information for all frequency input references
- configuration and status information for any frequency input reference
- overall information for all frequency output references
- configuration and status information for any frequency output reference
- configuration and status information for any frequency protection group

Step	Action

To display global synchronization information

1 Display global synchronization information:

```
sync show
```

To display the quality level hierarchy table

2 Display the quality level hierarchy table:

```
sync show ql-hierarchy-table
```

To display overall information for all frequency input references

3 Display overall information for all frequency input references:

```
sync input freq-ref show
```

To display configuration and status information for any frequency input reference

4 Display configuration and status information for any frequency input reference:

```
sync input freq-ref show freq-ref <STRING[32]>
```

where

freq-ref <STRING[32]> is the frequency input reference that you want to display.

To display overall information for all frequency output references

5 Display overall information for all frequency output references:

```
sync output freq-ref show
```


Copyright© 2012-2013 Ciena® Corporation

To display configuration and status information for any frequency output reference

6 Display configuration and status information for any frequency output reference:

```
sync output freq-ref show freq-ref <STRING[32]>
```

where

freq-ref <STRING[32]> is the frequency output reference that you want to display.

To display configuration and status information for any frequency protection group

7 Display configuration and status information for any frequency protection group:

```
sync protection-group show freq-group <STRING[32]>
```

where

freq-group <STRING[32]> is the frequency protection group that you want to display.

—end—

Example

The following example shows sample output for the sync show command.

```
sync show

+------------------------- SYNC GLOBAL CONFIGURATION --------------------+
| Parameter              | Value                                         |
+------------------------+-----------------------------------------------+
| Admin State            | Enabled                                       |
| Option Type            | Option 1                                      |
| Reversion Mode         | Revertive (default)                           |
| Wait-to-restore timer  | 10 minutes                                    |
+------------------------+-----------------------------------------------+
```

The following example shows sample output for the sync show ql-hierarchy-table command.

```
sync show ql-hierarchy-table

+-------- HIERARCHY OF QUALITY LEVELS -------+
|   QL    | Option I   | Option II |   Order   |
+---------+------------+-----------+-----------+
|     1   |            | QL-PRS    |  Highest  |
|     0   |            | QL-STU    |           |
|     2   | QL-PRC     |           |           |
|     7   |            | QL-ST2    |           |
|     4   | QL-SSU-A   | QL-TNC    |           |
|     8   | QL-SSU-B   |           |           |
|    13   |            | QL-ST3E   *           |
|    10   |            | QL-ST3/   |           |
|         |            | QL-EEC2   |           |
|    11   | QL-SEC/   *|           |           |
|         | QL-EEC1    |           |           |
```


Copyright© 2012-2013 Ciena® Corporation

```
|   12   |            | QL-SMC     |      |      |
|   14   |            | QL-PROV    |      |      |
|   15   | QL-DNU     | QL-DUS     |  Lowest   |
+--------+------------+------------+-----------+
* Local OCXO QL
```

The following example shows sample output for the sync input freq-ref show command.

```
sync input freq-ref show
```

```
+------------------------ SYNC INPUT FREQUENCY REFERENCES -------------------------+
|                            |  Timing  |       |Ov|              |                 |
|       Reference Name       |Interface |Proto  |Pr|   Oper QL    |  Operational Status  |
+----------------------------+----------+-------+--+--------------+-----------------+
|myBITSref                   |BITS-A    |BITS   | 1| 2 (QL-PRC)   |Selected         |
|port2                       |port 2    |SyncE  | 7|15 (QL-DNU)   |QL below threshold  |
+----------------------------+----------+-------+--+--------------+-----------------+
```

The following example shows sample output for the sync input freq-ref show freq-ref command.

```
sync input freq-ref show freq-ref
```

```
+------------------------ SYNC INPUT FREQUENCY REFERENCE --------------------+
| Parameter                | Value                                           |
+--------------------------+-------------------------------------------------+
| Reference Name           | myBITSref                                       |
| Timing Interface         | BITS-A                                          |
| Protection Group         | myPGroup1                                       |
| Operational Status       | Selected                                        |
+--------------------------+-------------------------------------------------+
| Protocol                 | BITS                                            |
| Mode                     | E1                                              |
| Format                   | ESF                                             |
| Encoding                 | AMI                                             |
| SSM Location             | SA4                                             |
+--------------------------+-------------------------------------------------+
| Override Priority        | 1                                               |
| Operational QL           | 2 (QL-PRC)                                      |
| Forced QL                | 2 (QL-PRC)                                      |
| RX QL                    | -                                               |
| QL Receive               | Disabled                                        |
+--------------------------+-------------------------------------------------+
```

The following example shows sample output for the sync output freq-ref show command.

```
sync output freq-ref show
```

```
+------------------------ SYNC OUTPUT FREQUENCY REFERENCES ------------------------+
|                            |  Timing  |       |              |                  |
|        Reference Name      |Interface |Proto  |   Oper QL    |  Operational Status  |
+----------------------------+----------+-------+--------------+------------------+
|synceGig10_out              |port 10   |SyncE  | 2 (QL-PRC)   |Active            |
|synceGig2_out               |port 7    |SyncE  |      -       |Loss of Signal    |
```


Copyright© 2012-2013 Ciena® Corporation

```
+---------------------------+--------+-----+------------+--------------------+
```

The following example shows sample output for the sync output freq-ref show freq-ref command.

```
sync output freq-ref show freq-ref synceGig10_out

+---------------------- SYNC OUTPUT FREQUENCY REFERENCE --------------------+
| Parameter             | Value                                             |
+-----------------------+---------------------------------------------------+
| Reference Name        | synceGig10_out                                    |
| Timing Interface      | port 10                                           |
| Operational Status    | Active                                            |
+-----------------------+---------------------------------------------------+
| Protocol              | SyncE                                             |
+-----------------------+---------------------------------------------------+
| TX QL                 | 2 (QL-PRC)                                        |
+-----------------------+---------------------------------------------------+
```

The following example shows sample output for the sync protection-group show freq-group command.

```
sync protection-group show freq-group myPGroup1

+------------------------ SYNC PROTECTION GROUP STATE ----------------------+
| Parameter               | Value                                           |
+-------------------------+-------------------------------------------------+
| Group Name              | group1                                          |
| Preferred Reference     | -                                               |
| Threshold QL            | 11 (QL-SEC)                                     |
| Admin State             | Enabled                                         |
| Operational Status      | Locked                                          |
| Duration                | 0:19:16:05                                      |
| Ref Switch Count        | 1                                               |
+-------------------------+-------------------------------------------------+

+----------------- SYNC PROTECTION GROUP FREQUENCY REFERENCES --------------+
|                       |  Timing |      |Ov|              |                  |
| Reference Name        |Interface|Proto |Pr|   Oper QL    | Operational Status|
+-----------------------+---------+------+--+--------------+------------------+
|Input2.1               |port 2.1 |SyncE| -|11 (QL-SEC)   |Selected          |
+-----------------------+---------+------+--+--------------+------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Copyright© 2012-2013 Ciena® Corporation

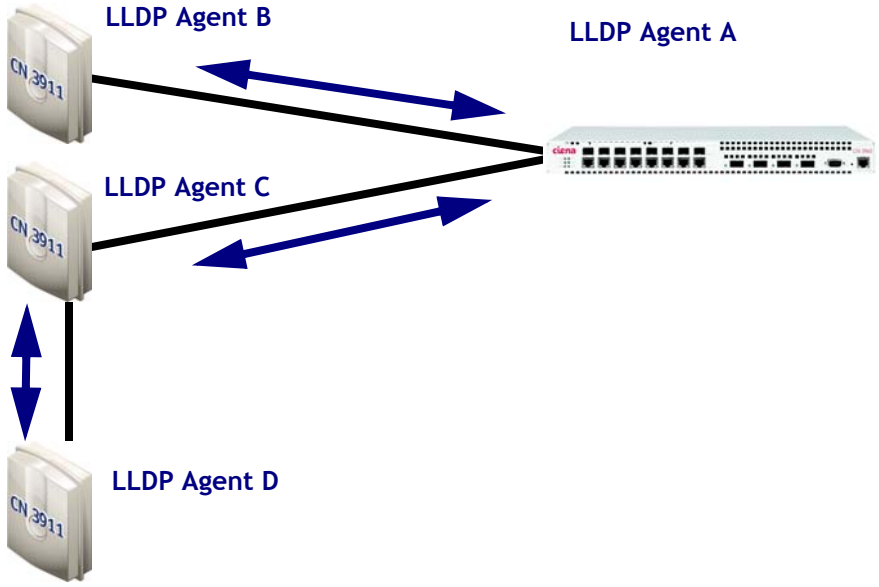
Link Layer Discovery Protocol (LLDP) configuration

The system software supports Link Layer Discovery Protocol (LLDP) as specified in the IEEE 802.1AB-2005 standard. Like Rapid Spanning Tree Protocol (RSTP), LLDP is a link-constrained, layer 2 protocol. This means that the exchange of LLDP messages only takes place between adjacent LLDP agents (devices) on the network, unless control frame tunneling is used to tunnel the LLDP PDUs through another device. LLDP agents communicate basic management information with each other by exchanging Link Layer Discovery Protocol Data Units (LLDPDUs). The information about a device, and its immediate neighbors is then retrievable through the standard LLDP MIBs using SNMP. It is important to note that LLDP operates only on physical ports.

Copyright© 2012-2013 Ciena® Corporation

As illustrated in Figure 6-1, LLDP Agent A exchanges LLDPDUs with LLDP Agent B and Agent C, but not with Agent D since it is not directly connected to Agent A.

Figure 6-1
Example of Architectural Relationship Between LLDP Agents



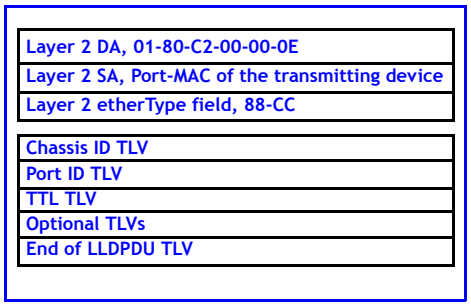
Note: Link Layer Discovery Protocol does not configure or control any traffic or devices on the network. Its primary role is to report discovered information to higher-layer management tools. It is not intended to act as a configuration protocol for remote systems nor as a mechanism to signal control information between ports.

Copyright© 2012-2013 Ciena® Corporation

LLDP TLVs

The LLDPDU is a layer 2 packet that consists of an L2 source, destination, and an Ethertype field, and 4 or more Type Length and Value fields (TLVs). The LLDP standard specifies 9 common TLV fields, 4 of which are mandatory TLVs and the remaining 5 are optional TLVs to carry information for broadcasting sender information.

Figure 6-2
LLDPDU Packet



Periodic LLDPDUs are sent out at a user-defined interval; however, LLDPDUs are also sent whenever LLDP TLV data has changed. An SNMP notification is then generated after changed data is received from a neighboring device. Upon receipt of this notification the SNMP management application polls the LLDP MIB objects to determine what information has changed.

Common TLVs

Table 6-1 lists LLDP TLVs and provides a description for each TLV. The system software supports the nine common TLVs specified by IEEE 802.1AB-2005.

Table 6-1
Common TLVs

TLV	Description
Common	
Chassis ID TLV	Identifies the chassis. This TLV contains the 802 LAN station associated with the transmitting LLDP agent and supports the “MAC Address” chassis ID subtype. This TLV is mandatory.
Port ID TLV	Identifies the port component that transmits the TLV. It supports the “Interface Alias” port ID subtype. This TLV is mandatory.
Time To Live (TTL) TLV	Indicates the number of seconds that the recipient LLDP agent is to regard the information associated with this LLDPDU to be valid. This TLV is mandatory.
Port Description TLV	Specifies the description for the port as an alphanumeric string. This TLV is optional.
System Name TLV	Specifies the assigned name as an alphanumeric string. This TLV is optional.
System Description TLV	Indicates the system description as an alphanumeric string. This TLV is optional.
System Capabilities TLV	Identifies the primary functions of the system and whether these primary functions are enabled. It supports the bridge capability. This TLV is optional.
Management Address TLV(s)	Identifies the IPv4 and IPv6 address associated with the local LLDP agent that can be used to reach higher layer entities to assist discovery by network management. This TLV is optional.
End of LLDPDU TLV	Defines the end of LLDPDU with all 0 values in two octets. This TLV is mandatory.

Copyright© 2012-2013 Ciena® Corporation

Table 6-1
Common TLVs

TLV	Description
802.1 Organizationally Specific TLVs	
Port VLAN ID TLV	Allows a port to advertise its VLAN ID (PVID) that is associated with untagged or priority tagged frames.
Port and Protocol VLAN ID TLV	The system software stores values received from the remote partner for this TLV but does not transmit this TLV.
VLAN Name TLV	The system software stores values received from the remote partner for this TLV but does not transmit this TLV.
Protocol ID TLV	Allows an 802 LAN station to advertise particular protocols that are accessible through the port. Currently, the system software advertises the following protocols: RSTP and 802.3ah OAM.
802.3 Organizationally Specific TLVs	
MAC/PHY Configuration/ Status TLV	Advertises autonegotiation support and status, PMD autonegotiation capability, and operational MAU type.
Power via MDI TLV	The system software stores values received from the remote partner for this TLV but does not transmit this TLV.
Link Aggregation TLV	Advertises whether the link is capable of being aggregated, whether the link is currently in an aggregation, and the port ID of the aggregation if it is in an aggregation. LLDP advertises Link Agg TLV in the PDU but Link Aggregation is not currently supported and is disabled.
Maximum Frame Size TLV	Advertises the maximum frame size capability.
Ciena Organizationally Specific TLVs	
Unknown TLVs	The system software stores all unknown TLVs for all valid LLDP PDUs for retrieval.
LLDP Graceful Shut Down	When an LLDP agent is administratively disabled, it executes a graceful shutdown handshake. It sends out the last LLDP PDU that specifies the TTL as zero in TTL TLV. When its counterpart receives this LLDPDU, it cleans up all the LLDP information received.

Copyright© 2012-2013 Ciena® Corporation

Feature Benefits

Using LLDP for network topology discovery offers the following benefits to the network provider:

- Accurate network topology discovery and management
- Support of standard tools using SNMP
- Multi-vendor interoperability.

Accurate Network Topology Discovery and Management

Many network management tools use layer 3 protocols to automate the discovery process and track topology configurations and changes. The use of layer 2 LLDP allows network management tools to quickly and accurately discover current network topologies and to track both intentional and unintentional topology changes. Network administrators and technicians can use the accurate and up-to-date topology information to more quickly diagnose network problems in the field.

Support of Standard Tools

Each LLDP agent maintains its own per-port table of information in a standard SNMP MIB. Updates are sent as needed to the closest connected device on the network. Users access the device and manage LLDP information through the command line interface (CLI). Users can display general information, for example, Chassis ID and port configuration. Users with superuser security privileges can configure LLDP options and manipulate TLVs. For example, a specific port could be configured to only transmit or receive LLDPDUs.

LLDPDUs are used to update standard LLDP Management information Databases (MIBs) allowing any standard SNMP application to monitor the information as it changes.

Multi-vendor Interoperability

LLDP is a standards-based protocol. Using LLDP rather than a proprietary topology discovery protocol allows the network provider to interoperate with non-Ciena devices that also support LLDP.

This chapter provides the following procedures for LLDP:

- “Configuring LLDP” on page 6-7
- “Configuring TLV transmission” on page 6-9
- “Displaying LLDP neighbors” on page 6-11
- “Enabling and disabling SNMP notifications” on page 6-12

Copyright© 2012-2013 Ciena® Corporation

Procedure 6-1
Configuring LLDP

LLDP is enabled by default on all ports, but can be enabled or disabled on a per-port basis.



CAUTION
Performance During Topology Discovery
The default settings for LLDP should be sufficient to ensure proper topology discovery in most networks. Since Ethernet Services Manager (ESM) topology discovery relies on LLDP for higher performance topology discovery, care should be taken when modifying the LLDP configuration. In certain topologies, LLDP does not need to be forwarded, for example, when using non-LLDP devices such as hubs.

Step	Action

1 Display the current state of the port.

```
lldp show port <port>
```

where

<port> is the port to be disabled.

2 Disable the port.

```
lldp set port <port> disable
```

where

<port> is the port to be disabled.

3 To verify the configuration of the port, execute the following command:

```
lldp show port <port> configuration
```

where

<port> is the port to be disabled.

—end—

Example

In the following configuration example, the user displays the state of port 10, then disables LLDP on this port. It should be noted as well that by using the `lldp show port` command, information about the remote port (neighbor) can be seen in the LLDP Remote Port TLV section of the output. This displays the MAC address, port number, and other information of the connected remote port.

Copyright© 2012-2013 Ciena® Corporation

Display the current state of port 10.

```
lldp show port 10
+----------- LLDP Port Configuration -----------+
|       |         |  Basic   |  Org Specific  |       |
|       | Admin   |TLV  Type | Dot1   |Dot3|   |       |
| Port  | State   |123456789|1234567|1234|   |Notif|
+------+-------+---------+-------+----+---+-----+
|10    | Tx-Rx |111111111|1  1111|1111|   | Off |
+------+-------+---------+-------+----+---+-----+
```

Set port 10 to disable.

```
lldp set port 10 disable
```

Verify the configuration of the port.

```
lldp show port 10 configuration
+----------- LLDP Port Configuration -----------+
|       |         |  Basic   |  Org Specific  |       |
|       | Admin   |TLV  Type | Dot1   |Dot3|   |       |
| Port  | State   |123456789|1234567|1234|   |Notif|
+------+-------+---------+-------+----+---+-----+
|10    |Disable|111111111|1  1111|1111|   | Off |
+------+-------+---------+-------+----+---+-----+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 6-2
Configuring TLV transmission

Some optional TLVs can be excluded from transmission in the LLDP PDU on a per-port basis.

Step	Action

1 Set TLV transmission parameters:

```
lldp tlvtx set port <port> mgmt-addr-local <on|off> mgmt-
addr-remote <on|off>  port-descr <on|off>  system-cap
<on|off>  system-descr <on|off>  system-name <on|off>
```

where

port <port>	is the port list.
mgmt-addr-local <on|off>	indicates whether to transmit the Local Management Address. The default value is on.
mgmt-addr-remote <on|off>	indicates whether to transmit the Remote Management Address. The default value is on.
port-descr <on|off>	indicates whether to transmit the port description. The default value is on.
system-cap <on|off>	indicates whether to transmit the system capabilities. The default value is on.
system-descr <on|off>	indicates whether to transmit the system description. The default value is on.
system-name <on|off>	indicates whether to transmit the system name. The default value is on.

2 Set TLV transmission dot1 parameters:

```
lldp tlvtx-dot1 set port <port> port-vlan-id <on|off>
protocol-id-dot1x <on|off> protocol-id-lacp <on|off>
protocol-id-oam <on|off>  protocol-id-stp <on|off>
```

where

port <port>	is the port list.
port-vlan-id <on|off>	indicates whether to transmit the port VLAN ID. The default value is on.
protocol-id-dot1x <on|off>	indicates whether to transmit the protocol ID xdot1x. The default value is on.

Copyright© 2012-2013 Ciena® Corporation

where

protocol-id-lacp <on|off>	indicates whether to transmit the protocol ID LACP. The default value is on.
protocol-id-oam <on|off>	indicates whether to transmit the protocol ID OAM. The default value is on.
protocol-id-stp <on|off>	indicates whether to transmit the protocol ID STP. The default value is on.

3 Set TLV transmission dot3 parameters:

```
lldp tlvtx-dot3 set port <port> link-agg <on|off> mac-
phy-config <on|off> max-frame-size <on|off> power-via-mdi
<on|off>
```

where

port <port>	is the port list.
link-agg <on|off>	indicates whether to transmit the link aggregation status. The default value is on.
mac-phy-config <on|off>	indicates whether to transmit the MAC Phy configuration. The default value is on.
max-frame-size <on|off>	indicates whether to transmit the maximum frame size. The default value is on.
power-via-mdi <on|off>	indicates whether to transmit the power via MDI status. The default value is on.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 6-3
Displaying LLDP neighbors

Display LLDP neighbors when you want to view the following information about neighboring devices:

- local port
- remote port
- remote management address
- chassis identification
- system name

Use the information to troubleshoot connectivity issues to other network elements.

Step	Action
1	Display neighboring devices: `lldp show neighbors`

—end—

Example

The following example shows sample output for the lldp show neighbors command.

```
lldp show neighbors
+------------------------------------------------------------------------+
|                                  LLDP Neighbors                        |
+------------------------------------------------------------------------+
| Remote Addr: 1.2.3.164                                                 |
|  Local Addr: 10.5.5.164                                                |
|  Local Addr: 2001:0db8:f018:1:202:5aff:fe01:b43a                       |
| Remote Addr: fe80::202:5aff:fe01:b449                                  |
| System Name: 3930-164                                                  |
| System Desc: 3930 Service Delivery Switch                              |
+------------+-----------------------------------------------------------+
|    Local   |                                Remote                     |
+------------+----------+------------------------------------------------+
|    Port    |    Port  |                        Info                    |
+------------+----------+------------------------------------------------+
|      5     |     5    |  Chassis Id: 00025A01BCD0                      |
|            |          |   Mgmt Addr: 10.5.5.178                        |
|            |          |              2001:0db8:f018:1:202:5aff:fe01:bcdf |
|            |          | System Name: 3916-178                          |
|            |          | System Desc: 3916 Service Delivery Switch      |
+------------+----------+------------------------------------------------+
...
+------------+----------+------------------------------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 6-4
Enabling and disabling SNMP notifications

LLDP supports the standard lldpRemTablesChange SNMP notification per port. This notification is disabled by default. When enabled, an SNMP notification is generated upon LLDP activity for the port when the value of the lldpStatsRemTableLastChangeTime changes.

You can

- enable SNMP notifications
- disable SNMP notifications

Step	Action

1 Enable notification of LLDP activity per port:

```
lldp set port <port> notification on
```

where

port <port>	is the physical port to enable notification of LLDP activity for.
notification on	enables notification.

2 Disable notification of LLDP activity per port:

```
lldp set port <port> notification off
```

where

port <port>	is the physical port to disable notification of LLDP activity for.
notification off	disables notification.

—end—

Copyright© 2012-2013 Ciena® Corporation

VLAN management

This chapter details the function of VLANs. It also details how to configure port settings to switch traffic properly through the network.

The following sections are in this chapter:

- "VLANs and traffic flow" on page 7-1
- "Behavior summary" on page 7-4
- "VLAN/port configuration" on page 7-4
- "VLAN translation" on page 7-5

Note: VLAN configuration requires installation of the Advanced Ethernet (AE) license. If you do not have an AE license, contact Ciena Sales.

VLANs and traffic flow

VLANs determine how traffic is forwarded within a device. The ingress and egress behavior of each port must be taken into account. The basic rules to govern traffic forwarded through the device are as follows:

- Port VLAN Membership
- Acceptable Frame Types (ingress rule)
- Port VLAN ID (PVID) (ingress rule)
- Ingress Filtering (ingress rule)
- Ingress Untagged Frames (ingress rule)
- Egress Untagged VLAN (egress rule)

Acceptable Frame Types

When a frame arrives at a port on a device with SAOS operating system, it is first examined to determine whether it is a tagged (contains a VLAN ID) or untagged frame. This is referred to as the frame type. Each port of the device has an associated Acceptable Frame Types (AFT) parameter that controls whether or not the port will accept untagged frames. A port can be configured to admit ONLY tagged frames or to admit ALL Frames.

Copyright© 2012-2013 Ciena® Corporation

When the value of the Acceptable Frame Types parameter is set to “VLAN tagged-only,” tagged frames are allowed to ingress while untagged or priority-tagged frames (i.e., a frame with no tag header, or a frame with a tag header that carries the null VLAN ID) received on the port are discarded. Tagged frames are then further processed according to the additional ingress rules set on the port. In contrast, if the Acceptable Frame Types parameter is set to “all”, all frames are allowed to ingress, regardless of their tag status.

Port VLAN ID

The full VLAN address range from 1 to 4094 can be supported per port. For untagged or priority-tagged frames, the next ingress rule applied is the Port VLAN ID (PVID). Each port on the device has an associated PVID. By default, the PVID value is the default VLAN ID (VLAN 1). When an untagged frame arrives at an ingress port (assuming that the port is set to accept all frame types) it may be forwarded to all ports that are members of the configured PVID (depending on additional ingress parameters set on the device). If an ingress frame is not tagged with the VLAN specified by the port’s PVID, then the setting on the port's Ingress Filter will apply.

Similarly, if a port is removed from a VLAN that is also its PVID, the PVID value remains the same and the port's Ingress Filter setting is therefore applied (refer to the VLAN Ingress Filter section).

The PVID assignment for aggregated ports is slightly different. When a physical port is added to an aggregation, it inherits the PVID setting of the aggregation group. However, when the port is removed from the aggregation, its own PVID setting becomes valid again.

> ***Note:*** PVIDs can only be set to the VLAN ID of configured VLANs, therefore the VLAN must exist before the corresponding PVID can be assigned. The port does not however, have to be a member of the VLAN in order to assign the PVID.

VLAN Ingress Filter

The last parameter applied to ingressing frames is the VLAN Ingress Filter. The VLAN Ingress Filter forwards or discards a frame based on the VLAN tag. When VLAN Ingress Filtering is enabled on a port (VLAN Ingress Filtering is enabled by default), an ingressing frame that is tagged with a VLAN to which the ingress port does not belong, will be discarded. For example, assuming that VLAN 200 is configured on the device, if the ingressing frame is tagged with VLAN 200, but the port is not a member of VLAN 200, the frame is discarded.

If VLAN Ingress Filtering is disabled, ports do not consider the VLAN membership of ingressing frames. In the above example, with VLAN Ingress Filtering disabled, the frame tagged with VLAN 200 may be forwarded to all

Copyright© 2012-2013 Ciena® Corporation

port members of VLAN 200, regardless of whether the ingress port belongs to VLAN 200. However, if VLAN 200 has not been configured on the device, the frame is then dropped.

Egress Untagged VLAN

Once a frame reaches its egress port on the device, one final VLAN operation is performed. The Egress Untagged VLAN (EUV) parameter determines whether the egressing frame will be forwarded tagged or untagged. As a frame egresses the device, its VLAN ID is compared to the EUV. If the VLAN ID is the same as the EUV, then the VLAN tag is stripped from the egressing frame and it is sent untagged. If the VLAN ID is not the same as the port's EUV, then the frame is sent out with its tag intact. By default, the EUV is VLAN 1, as is the PVID. When the PVID is changed, the EUV value automatically updates to match the PVID, until it is explicitly set.

For example, the default PVID and EUV are set to 1. If you set port's PVID to VLAN 200, and the EUV is automatically set to VLAN 200. A frame that is tagged with VLAN 200 would have its tag stripped as it egresses. In contrast a frame with a VLAN tag of 300 will egress with its tag intact.

However, if you set the EUV to 300, the PVID remains at 200. Now, the frame tagged with VLAN 200 retains its tag on egress, and the frame with a VLAN tag of 300 has its tag removed.

Copyright© 2012-2013 Ciena® Corporation

Behavior summary

The following tables summarize packet behavior.

Table 7-1
Ingress Behavior

Ingress Frame	AFT	PVID	Ingress Filter	Behavior
Tagged	Tagged-only	N/A	Disabled	The frame is forwarded if the VLAN exists on the device.
Tagged	Tagged-only	N/A	Enabled	The frame is forwarded if the ingress port is a member of the ingressing frame's VLAN.
Tagged	All	N/A	Disabled	The frame is forwarded if the VLAN exists on the device.
Tagged	All	N/A	Enabled	The frame is forwarded if the ingress port is a member of the ingressing frame's VLAN.
Untagged	Tagged-only	N/A	N/A	The frame is dropped.
Untagged	All	Any	Disabled	The frame is forwarded to all ports that are members of the ingressing port's PVID.
Untagged	All	Any	Enabled	The frame is forwarded to PVID members only if the ingress port is a member of the PVID.

Table 7-2
Egress Behavior

Frame State	EUV State	Egress Behavior
Tagged	VLAN ID = EUV	Frame sent untagged.
Tagged	VLAN ID ≠ EUV	Frame sent tagged.

VLAN/port configuration

All of the frame forwarding behaviors discussed in this chapter are port-based features, but require VLANs to tie them together. The VLAN can be thought of as an imaginary wire that connects the ingress port to the egress port(s). VLANs must be created using the CLI, MIB, or the Device Manager, prior to configuring other features on the device. A VLAN is identified by two basic parameters:

- VLAN ID, which is the value used to identify the VLAN
- VLAN Name, which is defined by the network operator. VLAN names may not begin with a number. This is an optional parameter.

Copyright© 2012-2013 Ciena® Corporation

VLAN translation

VLAN tags in a LAN often identify separate logical broadcast domains, for example, a workgroup. However, carriers are expected to support multiple enterprise networks on the same infrastructure and yet still preserve each LAN's VLAN schema. Manipulation of VLAN tags is therefore very useful.

For example, in a Multi-Tenant unit (MTU) environment, such as an office building with different subscribers, one switch may connect several subscriber networks in the building. Although each group is associated with a different physical port, care must be taken so that broadcast traffic is not exchanged between subscribers.

In an MTU environment, subscribers define their own VLAN schema, which may or may not be different from the carrier network or even other subscribers. VLAN translation can be used to remap or swap one VLAN ID with another VLAN ID.

VLAN translation is configured with port, VID, and VLAN entry which must exist as port members of the local VLAN. On ingress, VLAN translation looks up the frame's outer VID and maps the VID to the local VLAN. This local VLAN becomes the switching domain. On egress, VLAN translation looks up the local VLAN and marks the associated VID in the frame.

Note: MAC learning is performed on the local VLAN.

VLAN translation can not be configured for:

- A remote management VLAN
- A VLAN in use by an Ethernet Virtual Circuit
- A VLAN that is a member of a multicast service
- A reserved VLAN
- Ports that are virtual switch members

Table 7-3 shows the number of VLAN translation entries supported per platform.

Table 7-3
Supported VLAN Translation Entries

Platform	Entries Supported
3902, 3911, 3920, 3940	1000
3916, 3930, 3931, 3960, 5150	4096

Copyright© 2012-2013 Ciena® Corporation

Table 7-4 lists several features related to VLAN functionality and the support provided for VLAN translation.

Table 7-4
Feature Support in VLAN Translation

Feature	Support
Remote Management VLAN	Not supported.
PVID	Not supported.
Port untagged	Not supported.
Egress untagged VLAN	Not supported.
Broadcast Containment	Supported. However, classification is based on the raw frame VID, not the translated VID.
IGMP/Multicast services	Not supported.
CFM/Y.1731	Not supported.
Link aggregation	Supported on 3960 and 5150. Supported with a separate entry per physical port in the link aggregation group on 3911, 3920, 3940, and 5140 platforms. Not supported on 3902.
DHCP Relay	Supported. DHCP relay can be enabled on the local VLAN with VLAN translation entries. DHCP frames are edited according to the correct VIDs.

This chapter provides the following procedures for VLAN management:

- "Changing the TPID stamp for a VLAN" on page 7-7
- "Configuring a VLAN/port pair to ingress tagged traffic and egress tagged traffic" on page 7-8
- "Configuring a VLAN/port pair to ingress untagged traffic and egress untagged traffic" on page 7-10
- "Changing tag status" on page 7-11
- "Configuring hybrid traffic" on page 7-13
- "Emulating a tagged Ethernet port" on page 7-14
- "Translating a single NNI VLAN" on page 7-15
- "Translating a dual NNI VLAN" on page 7-17

Copyright© 2012-2013 Ciena® Corporation

Procedure 7-1
Changing the TPID stamp for a VLAN

By default, the outer Tag Protocol Identifier (TPID) of frames egressing a member port is stamped with a value of 8100. This setting can be changed per VLAN to 88a8 or 9100 or back to the default of 8100.

Step	Action

1 Change the TPID stamp for a VLAN:

```
vlan set vlan <Vlan> egress-tpid <8100|9100|88A8>
```

where

vlan <Vlan> is the VLAN.

2 Set the VLAN Ethernet policy to vlan-tpid for the port on the VLAN where the packets egress:

```
virtual-circuit ethernet set port <PortNameList> vlan-
ethertype-policy vlan-tpid
```

where

<PortNameList> is the port on the VLAN where the packages egress.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 7-2
Configuring a VLAN/port pair to ingress tagged traffic and egress tagged traffic

This scenario is the default behavior for the device. The only requirement is to add the ingress and egress ports to the same VLAN. Since the EUV and the PVID are set to VLAN 1 by default, tagged frames enter the switch and egress the switch with their original tag intact for the VLAN configured.

Note: By default, all ports are members of VLAN 1. Frames tagged with VLAN 1 are also forwarded since the PVID is set to VLAN 1 by default and is not changed in this example.

Step	Action

1 Create the VLAN.

```
vlan create vlan <VlanList> [name <String[31]>]
```

where

vlan <VlanList>	is the numeric vlan ID or IDs to create. You can specify one, a list, or a range.
name <String[31]>	is an optional name. This parameter applies if you are creating one vlan.

2 Add the ingress and egress ports to the VLAN.

```
vlan add vlan <VlanList> port <PortNameList>} [tag
{<value> | none}]
```

3 To verify the creation of the VLAN, enter the following command:

```
vlan show [port <PortNameList> | vlan <VlanList>]
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Example

The following example configures a port pair that receives frames tagged with VLAN 300 and egresses them with the original tag intact.

```
vlan create vlan 300
vlan add vlan 300 port 1,5
vlan show vlan 300
```

```
+---------------------- VLAN INFO -----------------------+
| Parameter             | Value                          |
+-----------------------+--------------------------------+
| VLAN ID               | 300                            |
| Name                  | VLAN#300                       |
| Features              |                                |
|-------------------------------------------------------+
|                      VLAN Members                     |
| Port      | VTag      | VS-Sub                         |
|-----------+-----------+--------------------------------+
| 1         | 300       | False                          |
| 5         | 300       | False                          |
|-----------+-----------+--------------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 7-3
Configuring a VLAN/port pair to ingress untagged traffic and egress untagged traffic

This scenario configures a port pair on a VLAN that receives untagged frames and egress them untagged:

Step	Action

1 Create the VLAN.

```
vlan create vlan <VlanList> [name <String[31]>]
```

where

vlan <VlanList> is the numeric vlan ID or IDs to create. You can specify one, a list, or a range.

name <String[31]> is an optional name. This parameter applies if you are creating one vlan.

2 Add the ingress and egress ports to the VLAN.

```
vlan add vlan <VlanList> port <PortNameList>
```

3 Set the AFT on all ports to All.

```
port set port <PortNameList> [acceptable-frame-type
<'all' | 'tagged-only'>]
```

4 Set the PVID on all ports to the VLAN used in Step 2. This command is used with ports not associated with a virtual switch. When the PVID is set, the egress untagged VLAN is automatically set to the same value as the PVID.

```
port set port <PortNameList> [pvid <VlanList>]
```

—end—

Example

The following example configures a VLAN port pair.

```
vlan create vlan 300
vlan add vlan 300 port 1,5
port set port 1,5 acceptable-frame-type all
port set port 1,5 pvid 300
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 7-4
Changing tag status

In this port pair scenario, traffic for the specified VLAN enters one port untagged and egresses its partner port tagged with a specified VLAN ID and 802.1.1d priority. When traffic flows in the opposite direction, frames tagged with the specified VLAN and 802.1d priority will egress untagged.

Step	Action

1 Create a VLAN.

```
vlan create vlan <VlanList> [name <String[31]>]
```

where

vlan <VlanList> is the numeric vlan ID or IDs to create. You can specify one, a list, or a range.

name <String[31]> is an optional name. This parameter applies if you are creating one vlan.

```
Add the ingress and egress ports to the VLAN.
vlan add vlan <VlanList> port <PortNameList>
```

2 Set the PVID on the ingress ports to the VLAN used in Step 2.

```
port set port <PortNameList> [pvid <VlanList>]
```

3 To set the ingress fixed .1d priority on the port, enter the following command:

```
port set port port <PortNameList> [fixed-rcos <NUMBER: 0-7>]
```

4 Set the AFT on all ports to All.

```
port set port <PortNameList> [acceptable-frame-type <'all' | 'tagged-only'>]
```

5 Set the EUV on the egress port (provider port) to a VLAN that is different than the VLAN used in Step 2.

```
port set port <PortNameList> [egress-untag-vlan <VlanList>]
```

—end—

Example

In our example, untagged frames ingressing on port 1 will egress port 5 tagged with a VLAN ID of 300 and 802.1d priority of 5. Frames that ingress on port 5 with a VLAN tag of 300 and 802.1d priority of 5 will egress port 1 untagged.

Copyright© 2012-2013 Ciena® Corporation

Assume now that port 1 is connected to the subscriber while port 5 connects to the provider network, the frame will pass between the subscriber and the provider networks, but the VLAN tag of 300 and 802.1d priority of 5, which may only have meaning in the provider network, will never be seen in the customer's network.

```
vlan create vlan 300
vlan add vlan 300 port 1,5
vlan create vlan 4094
vlan add vlan 4094 port 1,5
port set port 1,5 acceptable-frame-type all
port set port 1 pvid 300
port set port 1 fixed-rcos 5
port set port 5 egress-untag-vlan 4094
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 7-5
Configuring hybrid traffic

This scenario configures a port pair that will receive either tagged or untagged frames. Tagged frames will egress with their original tag intact and untagged frames will egress untagged.

Step	Action

1 Create a VLAN.

```
vlan create vlan <VlanList> [name <String[31]>]
```

where

vlan <VlanList> is the numeric vlan ID or IDs to create. You can specify one, a list, or a range.

name <String[31]> is an optional name. This parameter applies if you are creating one vlan.

2 Add the ingress and egress ports to the VLAN.

```
vlan add vlan <VlanList> port <PortNameList> [tag
{<value> | none}]
```

3 Set the AFT on the ingress ports to All.

```
port set port <PortNameList> [acceptable-frame-type
<'all' | 'tagged-only'>]
```

4 Set the PVID on the ingress ports to VLAN 1 (this is the default setting).

```
port set port <PortNameList> [pvid <VlanList>]
```

5 Set the EUV on the egress ports to VLAN 1 (this is the default setting).

```
port set port <PortNameList> [egress-untag-vlan
<VlanList>]
```

—end—

Example

In this example, any traffic tagged with VLAN 300 enters the switch tagged and leaves the switch still tagged with VLAN 300, while untagged traffic enters untagged and leave untagged.

```
vlan create vlan 300
vlan add vlan 300 port 1,5
port set port 1,5 acceptable-frame-type all
port set port 1 pvid 1
port set port 5 egress-untag-vlan 1
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 7-6
Emulating a tagged Ethernet port

Emulate the standard behavior of a tagged Ethernet port. Standard tagged Ethernet ports accept tagged frames and drop all untagged frames. To drop frames that are tagged with a VLAN to which the ingress port does not belong, do not disable VLAN Ingress Filtering.

Step	Action

1 Create a VLAN.

```
vlan create vlan <VlanList> [name <String[31]>]
```

where

vlan <VlanList>	is the numeric vlan ID or IDs to create. You can specify one, a list, or a range.
name <String[31]>	is an optional name. This parameter applies if you are creating one vlan.

2 Add the ingress and egress ports to the VLAN.

```
vlan add vlan <VlanList> port <PortNameList> [tag
{<value> | none}]
```

3 Set the AFT on the ports to Tagged-Only.

```
port set port <PortNameList> [acceptable-frame-type
<'all' | 'tagged-only'>]
```

4 Disable Ingress VLAN Filtering on the ports. Disable VLAN Ingress Filtering only if the ingress port is not a member of the tagged frame's VLAN.

```
port set <PortNameList> [vlan-ingress-filter <'on' |
'off'>]
```

5 Set the EUV on the ports to the VLAN used in Step 2

```
port set port <PortNameList> [egress-untag-vlan
<VlanList>]
```

—end—

Example

In the following example, VLAN filtering is disabled, allowing all tagged frames to ingress if the VLAN they are tagged with has been configured on the device. Frames that are tagged with a VLAN ID of 300 egress untagged.

```
vlan create vlan 300
vlan add vlan 300 port 1,5
port set port 1,5 acceptable-frame-type tagged-only
port set port 1,5 vlan-ingress-filter off
port set port 1,5 egress-untag-vlan 300
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 7-7
Translating a single NNI VLAN

Translate a single NNI VLAN.

Step	Action
1	Create a VLAN. `vlan create vlan <VlanList> [name <String[31]>]]` where vlan <VlanList> is the numeric vlan ID or IDs to create. You can specify one, a list, or a range. name <String[31]> is an optional name. This parameter applies if you are creating one vlan.
2	Add the ingress and egress ports to the VLAN. `vlan add vlan <VlanList> port <PortNameList> [tag {<value> | none}]`
3	Add the translation entry to map the local VLAN to the VID on egress. `vlan translate add port <PortNameList> vid <VID> vlan <VlanList>` where port <PortNameList> is the egress port vid <VlanList> vlan <VlanList> is the local VLAN
4	Verify the configuration. `vlan translate show`

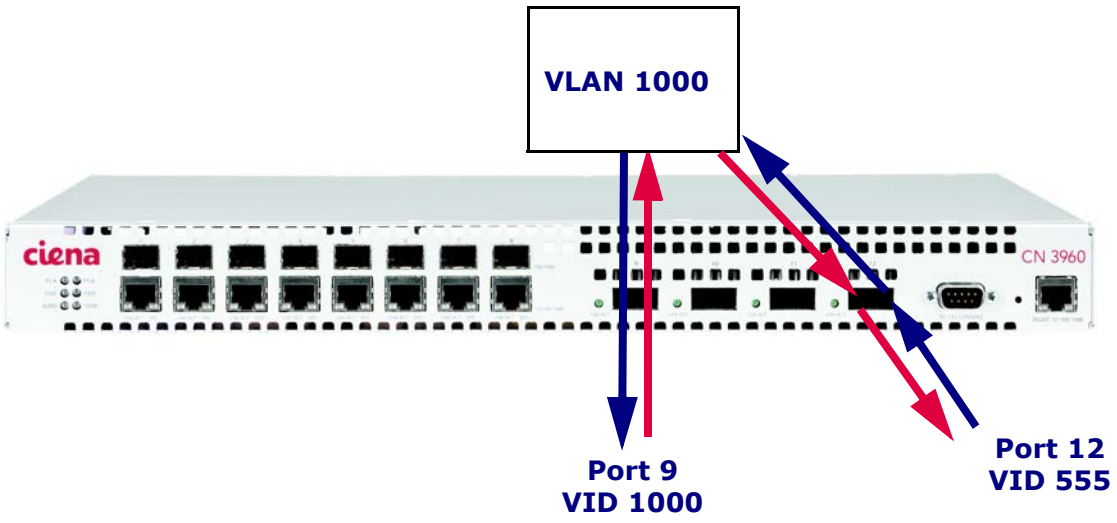
—end—

Example

In the configuration shown in Figure 7-1, Port 9 and Port 12 are members of local VLAN 1000. A VLAN translation entry maps Port 12 and VLAN 1000 to VID 555. When a packet arrives at Port 9 with VID 1000 it is mapped directly to local VLAN 1000, and then to VID 555. When a packet arrives at Port 12 with VID 555, it is mapped to local VLAN 1000.

Copyright© 2012-2013 Ciena® Corporation

Figure 7-1
Single NNI VLAN Translation



The following example creates the configuration.

```
vlan create vlan 1000
vlan add vlan 1000 port 9.12
vlan translate add port 12 vid 555 vlan 1000
vlan translate show

+-------------------------- VLAN TRANSLATE TABLE --------------------------+
|              |      Ingress                |        Egress               |           |
|              | Frame    | Swap-Internal    | Internal  | Swap-Frame      |           |
| Port         | VID      | VLAN             | VLAN      | VID             |           |
+--------------+----------+------------------+-----------+-----------------+-----------+
| 12           | 555      | 1000             | 1000      | 555             |           |
+--------------+----------+------------------+-----------+-----------------+-----------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 7-8
Translating a dual NNI VLAN

Translate a dual NNI VLAN.

Step	Action
1	Create a VLAN. `vlan create vlan <VlanList> [name <String[31]>]]` where vlan <VlanList> is the numeric vlan ID or IDs to create. You can specify one, a list, or a range. name <String[31]> is an optional name. This parameter applies if you are creating one vlan.
2	Add the ingress and egress ports to the VLAN. `vlan add vlan <VlanList> port <PortNameList> [tag {<value> | none}]`
3	Add the translation entry to map the local VLAN to the first VID on egress. `vlan translate add port <PortNameList> vid <VID> vlan <VlanList>` where port <PortNameList> is the egress port vid <VlanList> vlan <VlanList> is the local VLAN
4	Add the translation entry to map the local VLAN to the second VID on egress. `vlan translate add port <PortNameList> vid <VID> vlan <VlanList>` where port <PortNameList> is the egress port vid <VlanList> vlan <VlanList> is the local VLAN
5	Verify the configuration. `vlan translate show`

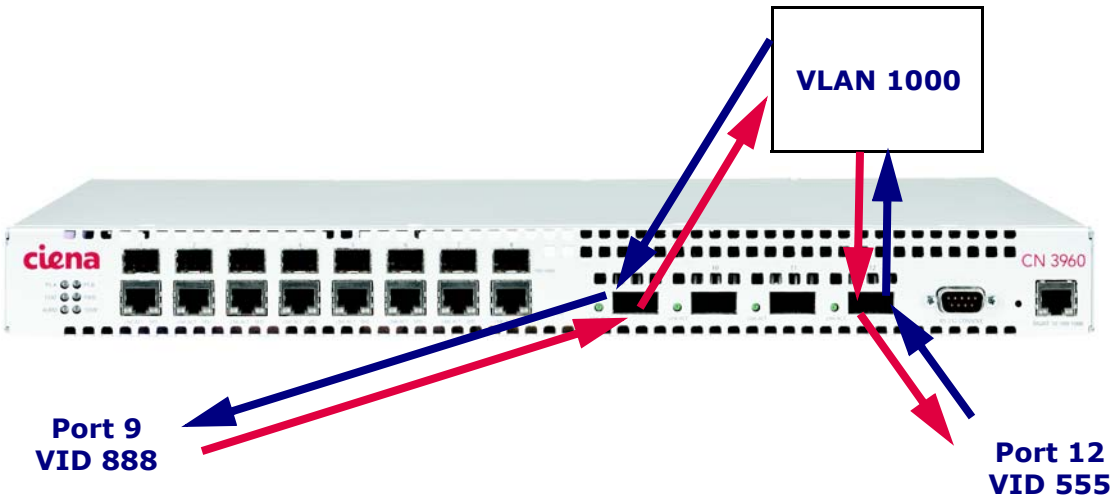
—end—

Copyright© 2012-2013 Ciena® Corporation

Example

In the configuration shown in Figure 7-2, Port 9 and Port 12 are members of local VLAN 1000. A VLAN translation entry maps Port 12 and VLAN 1000 to VID 555. Another maps Port 9 and VLAN 1000 to VID 888. When a packet arrives at Port 9 with VID 888 it is mapped to local VLAN 1000 and then to VID 555. Likewise, when a packet arrives at Port 12 with VID 555, it is mapped to local VLAN 1000 and then to VID 555.

Figure 7-2
Dual NNI VLAN Translation



The following example creates the configuration.

```
vlan create vlan 1000
vlan add vlan 1000 port 9,12
vlan translate add port 9 vid 888 vlan 1000
vlan translate add port 12 vid 555 vlan 1000
vlan translate show
```

```
+------------------------- VLAN TRANSLATE TABLE --------------------------+
|              |      Ingress                |        Egress               |           |
|              | Frame    | Swap-Internal    | Internal  | Swap-Frame      |           |
| Port         | VID      | VLAN             | VLAN      | VID             |           |
+--------------+----------+------------------+-----------+-----------------+-----------+
| 9            | 888      | 1000             | 1000      | 888             |           |
| 12           | 555      | 1000             | 1000      | 555             |           |
+--------------+----------+------------------+-----------+-----------------+-----------+
```


Copyright© 2012-2013 Ciena® Corporation

IP management

This chapter shows the commands for configuring static IP routing. 39xx/51xx switches support IPv4 and IPv6.

Note 1: To configure and run IP static routing, you need to install the Advanced-OAM license key. To obtain the Advanced-OAM license key, contact Ciena Sales.

Note 2: For an example of configuring Static IP routing to support RFC 2544 benchmark testing refer to 009-3203-009, SAOS 6.10 Fault and Performance Management.

Static IP routing comprises the following components:

- IP Interfaces, which enables IP packet sending and receiving.
- IP Loopback, which enables IP packets to ingress and egress over the same interface.
- Static IP routes, which defines static IP routes.
- Forward Information Base (FIB), which stores static route and forwarding information.
- Static Address Resolution Protocol routes, which defines ARP static routes.
- Adjacency Information Base (AIB), which stores ARP static route and adjacency information.

IPv6

IPv6 is the successor to IPv4. IPv6 provides a larger address space and greater flexibility when assigning addresses.

IPv6 address format

IPv6 addresses have the following characteristics:

- Addresses are 128-bits in length and are written as eight groups of 4 hexadecimal digits. For example:
  0001:0002:0003:0004:0005:0006:0007:0008
- Fields are separated by colons (:)

Copyright© 2012-2013 Ciena® Corporation

- Numeric field values are 0-FFFF
- Two colons together (::) can be used as a shorthand to represent some number of all-zeros hexadecimal fields. The doubled colon can appear at the beginning, the middle, or at the end of the address. Only one pair of doubled colons can be used in an address. For example: ::1 or f:: are both valid addresses.
- Port numbers are added to IPv6 addresses by enclosing the IP address in square brackets and then a colon followed by the port number. Because square brackets can be interpreted as regular expressions, the IP address and port number are enclosed in double quotes.

Ciena platforms accept most address forms called out in RFC4291, and emit only address forms called out in RFC5952, but (per this RFC section 4.2.2) Ciena platforms do not accept a double-colon that does not compress away more than 16 bits of address, that is, 1:2:3:4:5:6::8 is not legal.

IP address usage

On an IPv6-enabled network, each network element can have multiple multicast, anycast, and unicast addresses. IPv6 is most secure when used in a unicast deployment, that is, one host to another host. Multicast and anycast use one-to-many connections and are not recommended in a secure environment.

Copyright© 2012-2013 Ciena® Corporation

Table 8-1 summarizes IPv6 address types.

Table 8-1
IPv6 address types and IPv4 equivalent

IPv6 name	Description	IPv4 equivalent	Description
Anycast	One-to-nearest. Uses Global Unicast Addresses. Routers only. Discovery uses.	Unique protocols in IPv4, for example, IGMP.	Addresses are indistinguishable from a normal unicast address. Anycast (router-to-router) is also used with IPv4 addresses, specifically with DNS root servers, though there may be other instances.
Global Unicast	Globally unique. Fully routable. Assigned by IANA/ Regional Internet Registries (RIRs).	Global IP address	IPv6 and IPv4 similar but IPv6 can have other scoped addresses.
Link-Local	Local LAN only. Automatically assigned based on MAC. Cannot be routed outside local LAN. Scoped address concept new to IPv6. Multicast can also be scoped to link.	No real equivalent. Assigned IPv4 over ARP'd MAC.	Scoped address concept new to IPv6. Multicast may also be scoped to link-local (RFC 4489).
Loopback	Local interface scope.	Same as IPv4 127.0.0.1	Same function
Multicast	One-to-many. Hierarchy of multicasting.	Similar to IPv4 Class D	Significantly more powerful than IPv4 version. No broadcast in IPv6, replaced by multicast. Multicast may also be scoped to link-local (RFC 4489).
Site-Local	Optional. Local Site only. Cannot be routed over Internet. Assigned by user.	Private network address with multi-homed interface is closest equivalent.	Scoped address concept new to IPv6. Unlike the IPv4 private network address the IPv6 device can have, Link-Local, Site-Local and a Global Unicast address. Site-Local while continuing to exist in the IPv6 specification is currently not supported.

Copyright© 2012-2013 Ciena® Corporation

IPv6 architecture allows an interface to have multiple addresses, for example, link-local or global. 39xx/51xx switches also support IPv6 and IPv4 addresses. The rules for selecting the address to use in a particular situation are the following:

- Address pairs of the same scope or type (link-local, global) are preferred.
- A smaller scope for the Destination address is preferred, that is, use the smallest scope possible.
- A non-deprecated address is preferred.
- Transitional addresses, that is, IPv6 to IPv4 addresses, are not used if native IPv6 addresses are available.
- If all criteria are similar, address pairs with the longest common prefix are preferred.
- For source addresses, global addresses are preferred over temporary addresses.
- The attribute “interface” when adding a gateway is only required if the gateway address is an IPv6 link-local address.

Each interface can have multiple IPv6 addresses in addition to the IPv4 addresses. 39xx/51xx switches support the use of up to 16 IPv6 addresses per interface.

This chapter provides the following procedures for IP management:

- “Creating interfaces” on page 8-8
- “Deleting an IP or loopback interface” on page 8-10
- “Modifying an IP or loopback interface” on page 8-11
- “Disabling an IP interface” on page 8-12
- “Enabling an IP interface” on page 8-13
- “Configuring IPv6 interfaces manually” on page 8-14
- “Displaying an IP interface” on page 8-16
- “Adding an IP route” on page 8-20
- “Removing an IP route” on page 8-21
- “Displaying the routing table” on page 8-22
- “Displaying FIB entries” on page 8-23
- “Enabling Layer 3 switching” on page 8-25
- “Disabling Layer 3 switching” on page 8-26
- “Adding a static ARP entry” on page 8-31
- “Removing static ARP entries” on page 8-32

Copyright© 2012-2013 Ciena® Corporation

- "Displaying static ARP entries" on page 8-50
- "Displaying AIB information" on page 8-51

EVC ping

EVC ping is used by network operators to test connectivity through end-to-end networks through an Ethernet Virtual Connection (EVC). Network operators can initiate a ping of the remote EVC endpoint (NTE-2) before scheduling deployment and testing of new equipment. EVC ping is configured using IP interface and virtual switch members.

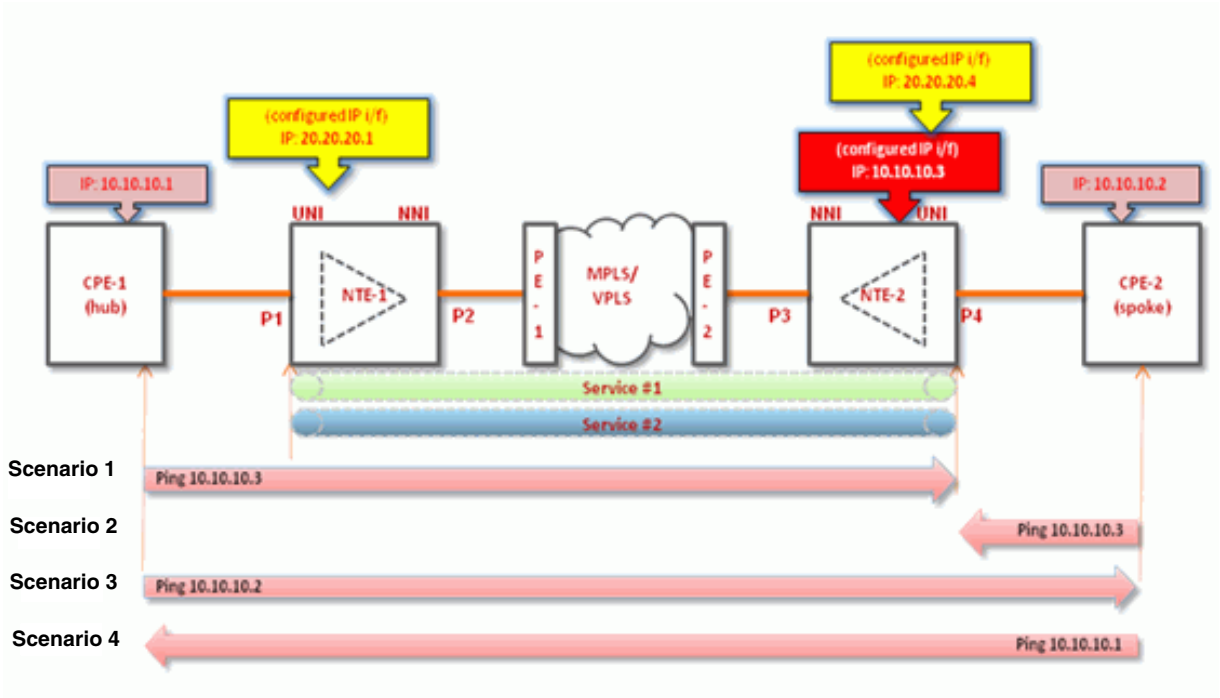
EVC ping can be used for EVPL and EPL services in IPv4-based networks.

When EVC ping is enabled, the NTE must be configured with an IP interface whose IP address is coordinated with the network operator. Each IP interface must be in its own subnet, and must not already be configured on the NTE.

Frames with any DSCP priority are accepted: frame treatment is subject to the provisioned SLAs in the network. No special priority is given to this ping traffic.

Figure 8-1 shows EVC ping in the end-to-end network.

Figure 8-1
EVC ping in the end-to-end network



Copyright© 2012-2013 Ciena® Corporation

Table 8-2 lists EVC ping use cases. Port-based service can have tagged or untagged traffic sent from the network operator. VLAN-based service must have tagged traffic, coordinated with the network operator.

Table 8-2
EVC ping use cases

Use case	Service on NTE-1 (hub)	Ping/ARP packet on NTE-1 UNI	Service on NTE-2 (spoke)	VLAN tag stack depth at NTE-2 NNI	Ping/ARP packet on NTE-2 UNI
A	port-based	tagged	port-based	2	tagged
B	port-based	untagged	port-based	1	untagged
C	VLAN-based	tagged	port-based	1	untagged
D	port-based	untagged	VLAN-based	2	tagged
E	VLAN-based	tagged	VLAN-based	2	tagged

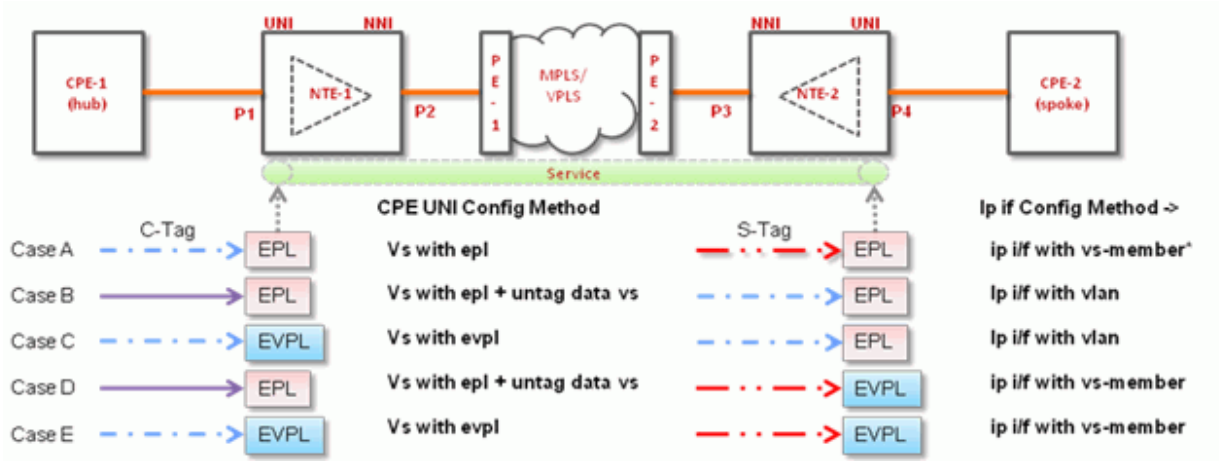
EVC ping can be configured for

- Use case A: virtual switch with EPL to IP interface with virtual switch member
- Use case B: virtual switch with EPL and untagged data virtual switch to IP interface with VLAN
- Use case C: virtual switch with EVPL to IP interface with VLAN
- Use case D: virtual switch with EPL and untagged data virtual switch to IP interface with virtual switch member
- Use case E: virtual switch with EVPL to IP interface with virtual switch member

Figure 8-2 shows EVC ping use cases.

Copyright© 2012-2013 Ciena® Corporation

Figure 8-2
EVC ping use cases



This chapter provides the following procedures for EVC ping:

- “Configuring EVC ping for use case A” on page 8-33
- “Configuring EVC ping for use case B” on page 8-36
- “Configuring EVC ping for use case C” on page 8-39
- “Configuring EVC ping for use case D” on page 8-42
- “Configuring EVC ping for use case E” on page 8-46

Copyright© 2012-2013 Ciena® Corporation

Procedure 8-1
Creating interfaces

You can create an:

- IP interface
- loopback interface

Create an IP interface to enable the sending and receiving of Layer 3 traffic.

Create a loopback interface when you need the Layer 3 traffic to ingress and egress from the same interface.

Note: The IP address and subnet mask for an IP interface cannot define a subnet that is already in use by a local or remote interface.

Step	Action

To create an IP interface

1 Create an IP interface.

```
interface create ip-interface <ip-interface> {ip <IP
address with mask>} [vlan <VLAN>] [mtu <NUMBER: 1500-
9216>] [mac <MAC address: XX:XX:XX:XX:XX:XX>]
[ip-forwarding <on|off>] [service-mac <benchmark | local-
mgmt | remote-mgmt>]
```

where

ip-interface <ip-interface>	is a 15-character string to identify the interface
ip <IP address with mask>	is the IPv4 address.
vlan <VLAN>	is the VLAN ID for this interface (port service). The VLAN is optional. When this occurs the ip-interface will obtain its L2 mapping through the use of virtual-switch member configuration.
[mtu <NUMBER: 1500-9216>]	is the interface MTU.

Copyright© 2012-2013 Ciena® Corporation

where

[mac <MAC address: XX:XX:XX:XX:XX:XX>]	is the MAC address.
ip-forwarding <on|off>	determines whether IP forwarding is on or off. The default value is off.
service-mac <benchmark | local-mgmt | remote-mgmt>]	is the service name associated with the MAC address. Alternate way of specifying the MAC address.

To create a loopback interface

2 Create a loopback interface:

```
interface create loopback <loopback> {ip <IP address>}
```

where

loopback <loopback>	is the loopback interface name.
ip <IP address>	is the interface IP address.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 8-2
Deleting an IP or loopback interface

Delete an IP or loopback interface when it is no longer required.

Step	Action
1	Delete an IP or loopback interface: `interface delete ip-interface <ip-interface>` where ip-interface <ip-interface> is the name of the interface to be deleted.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 8-3
Modifying an IP or loopback interface

Modify an IP or loopback interface to change the setting, for example, to increase the maximum frame size to allow bigger frames.

Step	Action

1 Modify an IP or loopback interface:

```
interface set ip-interface <Interface> {[ip <IpAddress>],
[ip-forwarding <on|off>], [mtu <1500-9216>], [subnet-mask
<SubnetMask>]}
```

where

<Interface>	is the IP or loopback interface to be modified.
<IpAddress>	is the IPv4 address.
ip forwarding	determines whether IP forwarding is on or off. The default value is off.
mtu	is the maximum transmission unit in bytes (maximum frame size). The default value is 1500.
<SubnetMask>	is the subnet mask in dotted decimal format.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 8-4
Disabling an IP interface

Disable an IP interface.

Step	Action
1	Disable an IP interface: `interface disable ip-interface <Interface>` where <Interface> is the IP interface to be disabled.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 8-5
Enabling an IP interface

Enable an IP Interface.

Step	Action
1	Enable an IP interface: `interface enable ip-interface <Interface>` where <Interface> is the IP interface to be enabled.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 8-6
Configuring IPv6 interfaces manually

Manually configure IPv6 interfaces.

Step	Action

To add the local interface IPv6 address and subnet mask

1 Add the local interface IPv6 address and subnet mask:

```
interface local add ip <ip>
```

where

ip <ip> is the IPv6 address of interface.

To remove the local interface IPv6 address and subnet mask

2 Remove the local interface IPv6 address and subnet mask:

```
interface local remove ip <ip>
```

where

ip <ip> is the IPv6 address of interface.

To add the remote interface IPv6 address and subnet mask

3 Add the remote interface IPv6 address and subnet mask:

```
interface remote add ip <ip>
```

where

ip <ip> is the IPv6 address of interface.

To remove the remote interface IPv6 address and subnet mask

4 Remove the remote interface IPv6 address and subnet mask:

```
interface remote remove ip <ip>
```

where

ip <ip> is the IPv6 address of interface.

To add an IPv6 gateway address

5 Add an IPv6 gateway address:

Copyright© 2012-2013 Ciena® Corporation

```
interface add { [gateway <gateway>] [interface
<local|remote>] }
```

where

gateway <gateway>	is the IPv6 gateway to be added. It is best practice to use the link-local address of the router (starting with FE80::) as a gateway rather than the global address of the router, as this allows ICMPv6 type 137 redirect messages to work correctly, thus enhancing network reachability.
interface <local|remote>	is the management interface.

To remove an IPv6 gateway address

6 Remove an IPv6 gateway address:

```
interface remove { [gateway <gateway>] [interface
<local|remote>] }
```

where

gateway <gateway>	is the IPv6 gateway to be added.
interface <local|remote>	is the management interface.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 8-7
Displaying an IP interface

Display an IP interface to verify the configuration.

The Interface show command shows only operational addresses. The Interface Management Table includes local, remote and IP interfaces. IP interfaces are only IPv4 supported. Local and remote interfaces support both IPv4 and IPv6.

The interface show command for local and remote interfaces displays the source of configuration of the specified IP address, as shown in Table 8-3.

Table 8-3
Source of configuration of the specified IP address

Source	Description
Manual	Specifies configuration by user
SLAAC	Configured with Stateless Auto configuration through Router Advertisement messages
DHCP	Configured through DHCP Server
Local	Local Addresses including Link Local Addresses

Step	Action

To display IP interfaces

1 Display IP interfaces:

```
interface show
```

To display local IP interfaces

2 Display local IP interfaces:

```
interface local show
```

To display remote IP interfaces

3 Display remote IP interfaces:

```
interface remote show
```


Copyright© 2012-2013 Ciena® Corporation

To display an IP interface

4 Display an IP interface:

```
interface show [ip-interface <Interface>]
```

where

<Interface> is the IP interface to be displayed.

—end—

Example

The following example shows sample output for the interface show command.

```
interface show
+---------------------------------- INTERFACE MANAGEMENT -----------------------------+
| Name                | Management         | IP Address/Prefix                          |
|                     | Domain             |                                            |
+---------------------+--------------------+--------------------------------------------+
| local               | VLAN 0             | 10.1.29.255/21                             |
| local               | VLAN 0             | fdbf:bc21:b16a:1123:2021:5aff:fe01:b910/64 |
| local               | VLAN 0             | fe80::202:5aff:fe01:b910/64                |
| remote              | VLAN 123           | 10.2.29.253/21                             |
| remote              | VLAN 123           | fe80::202:5aff:fe01:b91f/64                |
| fred                | VS Abcde67890abcde | 5.5.5.1/24                                 |
+---------------------+--------------------+--------------------------------------------+

+-------------- TCP/IP/STACK OPERATIONAL STATE ---------------+
| Parameter           | Value                                 |
+---------------------+---------------------------------------+
| IPv4 Gateway        | 10.1.24.1                             |
| IPv6 Gateway        | fe80::214:22ff:fe73:3987              |
| IPv4 Forwarding     | Off                                   |
| Default DSCP        | 0                                     |
+---------------------+---------------------------------------+
```

Note that fe80::202:5aff:fe01:b910/64 and fe80::202:5aff:fe01:b91f/64 are link-local addresses. Link-local addresses are partially derived from the chassis MAC address of the switch and are not routable outside of the local network.

The following example shows sample output for the interface local show command.

```
interface local show

+-------- INTERFACE OPERATIONAL STATE --------+
| Parameter           | Value                 |
+---------------------+-----------------------+
| Name                | local                 |
| Index               | 2                     |
| Admin State         | Enabled               |
| Oper State          | Enabled               |
```


Copyright© 2012-2013 Ciena® Corporation

```
| MAC Address          | 00:02:5a:01:b9:10     |
| Management Domain    | VLAN 0                |
| Priority             | 0                     |
| MTU                  | 1500                  |
+----------------------+-----------------------+

+------------------------- ADMIN INTERFACE ADDRESSES --------------------------+
| Parameter            | Value                                                 |
+----------------------+-------------------------------------------------------+
| IPv4 Addr/Mask       | 10.1.30.254/21                                        |
| IPv6 Addr/Mask       | 2001:1234:f00d:cafe:1111:1202:1001:1111/64            |
+----------------------+-------------------------------------------------------+

+---------------------- OPERATIONAL INTERFACE ADDRESSES ----------------------+
| Parameter            | Value                                       | Source   |
+----------------------+---------------------------------------------+----------+
| IPv4 Addr/Mask       | 10.1.30.7/21                                | DHCP     |
| IPv4 Broadcast Addr  | 10.1.31.255                                 |          |
| IPv6 Addr/Mask       | 2001:1234:f00d:cafe:1111:1202:1001:1111/64  | Manual   |
|                      | fe80::202:5aff:fe01:b990/64                 | Local    |
+----------------------+---------------------------------------------+----------+
```

The following example shows sample output for the interface remote show command.

```
interface remote show

+-------- INTERFACE OPERATIONAL STATE --------+
| Parameter            | Value                |
+----------------------+----------------------+
| Name                 | remote               |
| Index                | 15                   |
| Admin State          | Enabled              |
| Oper State           | Enabled              |
| MAC Address          | 00:02:5a:01:b9:1f    |
| Management Domain    | VLAN 123             |
| Priority             | 7                    |
| MTU                  | 1500                 |
+----------------------+----------------------+

+------------------------- ADMIN INTERFACE ADDRESSES --------------------------+
| Parameter            | Value                                                 |
+----------------------+-------------------------------------------------------+
| IPv4 Addr/Mask       | Not configured                                        |
| IPv6 Addr/Mask       | Not configured                                        |
+----------------------+-------------------------------------------------------+

+---------------------- OPERATIONAL INTERFACE ADDRESSES ----------------------+
| Parameter            | Value                                       | Source   |
+----------------------+---------------------------------------------+----------+
| IPv6 Addr/Mask       | fe80::202:5aff:fe01:b91f/64                 | Local    |
+----------------------+---------------------------------------------+----------+
```


Copyright© 2012-2013 Ciena® Corporation

The following example shows sample output for the show ip-interface command.

```
interface show ip-interface

+------------------------------ IP INTERFACE TABLE -----------------------------+
| Name                |  State  | Management          | MTU  | IP Address/Prefix  |
|                     |Adm |Oper| Domain              |      |                    |
+---------------------+----+----+---------------------+------+--------------------+
| joe                 |Ena |Dis | VLAN 100            | 1500 | 19.34.34.1/21      |
| fred                |Ena |Ena | VS Abcde67890abcde | 1500 | 5.5.5.1/24         |
+---------------------+----+----+---------------------+------+--------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 8-8
Adding an IP route

Add an IP route.

Step	Action

1 Add an IP route:

```
ip route add {destination <IpAddressWithMask>} {gateway
<IpAddress>} [metric <#>] [domain-name <routing-domain-
name>]
```

where

destination <IpAddressWithMask>	is the destination route IP address in CIDR notation.
gateway <IpAddress>	is the gateway IP address for the associated IP or loopback interface.
metric <#>	is the interface route cost metric. The default value is 0.
[domain-name <routing-domain-name>]	is the routing domain.

—end—

Example

The following example adds an IP route with a destination IP address of 5.5.5.0/24, a gateway address of 6.6.6.6, and an interface route cost metric of 10.

```
ip route add destination 5.5.5.0/24 gateway 6.6.6.6 metric 10
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 8-9
Removing an IP route

Remove an IP route when the IP route is no longer needed.

Step	Action

1 Remove an IP route:

```
ip route remove [destination <IpAddressWithMask>],
[gateway <IpAddress>], |[metric <#>] [domain-name
<routing-domain-name>]
```

where

destination <IpAddressWithMask>	is the destination route IP address in CIDR notation.
gateway <IpAddress>	is the gateway IP address for the associated IP or loopback interface.
metric <#>	is the interface route cost metric. The default value is 0.
[domain-name <routing-domain-name>]	is the routing domain.

—end—

Example

The following example removes a static IP route with a destination IP address of 5.5.5.0/24, a gateway address of 6.6.6.6, and an interface route cost metric of 10.

```
ip route remove destination 5.5.5.0/24 gateway 6.6.6.6 metric 10
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 8-10
Displaying the routing table

Display the routing table to verify configuration.

Step	Action
1	Display the routing table: `ip route show [domain-name <routing-domain-name>]` where [domain-name <routing-domain-name>] is the routing domain.

—end—

Example

The following example shows sample output for the ip route show command.

```
> ip route show

+------------------------------- ROUTING TABLE --------------------------------+
|Act| Destination     | Gateway         | Genmask         |Metric    |Intf|Prot|
+---+-----------------+-----------------+-----------------+----------+----+----+
|   | 10.0.0.0        | 1.1.1.1         | 255.0.0.0       | 0        | 5  | Inv|
+---+-----------------+-----------------+-----------------+----------+----+----+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 8-11
Displaying FIB entries

After creating a static IP route, the forwarding information is stored in the FIB. The summary view is displayed by default.

You can display:

- the FIB entry for the default route
- a summary of FIB entries
- details for each FIB entry

—end—

Step	Action

To display the FIB entry for the default route

1 Display the FIB entry for the default route:

```
ip fib show [destination <Ip Address with mask>] [default]
```

where

destination <Ip Address with mask>	is the destination route IP address/mask.
[default]	displays the default route, if any configured.

To display a summary of FIB entries

2 Display a summary of FIB entries:

```
ip fib show [destination <Ip Address with mask>][summary]
```

where

destination <Ip Address with mask>	is the destination route IP address/mask.
[summary]	displays summary.

Copyright© 2012-2013 Ciena® Corporation

To display details for each FIB entry

3 Display details for each FIB entry:

```
ip fib show [destination <Ip Address with Mask>][details]
[summary]
```

where

destination <Ip Address with Mask>	is the destination route IP address.
[details]	displays details.

—end—

Example

The following example shows sample output for the ip fib show command. The summary view is displayed by default.

```
> ip fib show

+------------------+----------------+---+---+----+------+-----------+
| Destination/     |  NexthopIp     |oif|Act|Path|Entry |Path Cost  |
|  NetMaskLen      |                |Idx|ion|Type|Id    |           |
+------------------+----------------+---+---+----+------+-----------+
|1.1.1.0/24        |-               |-  |loc|IF  |0     |0          |
|1.1.1.0/32        |-               |-  |loc|IF  |0     |0          |
|1.1.1.1/32        |-               |-  |loc|IF  |0     |0          |
|1.1.1.2/32        |1.1.1.2         |5  |loc|ADJ |0     |0          |
|1.1.1.255/32      |-               |-  |loc|IF  |0     |0          |
+------------------+----------------+---+---+----+------+-----------+
| oif: Outgoing Interface                                           |
| entryId: Internal Id from Route Module (=0 local)                 |
| Action: fwd: Forward, loc: Local, rej: Reject, bck: BlackHole     |
| C: Conn, LC: LocalConn, DC: DirectlyConn, Sta: Static             |
| PathType: I1I: Isis L1 Int, I2I: Isis L2 Int                      |
|  I1E: Isis L1 Ext, I2E: Isis L2 Ext, OIA: Ospf IntraArea          |
|  OA: Ospf InterArea, O1E: Ospf Type1 Ext, O2E: Ospf Type2 Ext     |
|  IC: Icmp, Oth: Other, IF: Local Interface, ADJ: Adjacency        |
|  unk: Unknown                                                     |
+-------------------------------------------------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 8-12
Enabling Layer 3 switching

FIB Layer 3 switching is enabled by default. In order for static IP routing to support layer 3 traffic, such as for RFC 2544 and MPLS, FIB Layer 3 switching must be enabled.

Step	Action
1	Enable Layer 3 switching: `ip fib l3-switching enable`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 8-13
Disabling Layer 3 switching

Disable Layer 3 switching when MPLS and benchmarking are no longer required.

Step	Action
1	Disable Layer 3 switching: `ip fib l3-switching disable`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 8-14
Displaying the status of Layer 3 switching

Display the status of Layer 3 switching to learn whether Layer 3 switching is enabled or disabled.

Step	Action
1	Display the status of Layer 3 switching: `ip fib l3-switching show`

—end—

Example

The followiing example shows sample output of the ip fib l3-switching show command.

```
> ip fib l3-switching show

L3-Switching Enabled
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 8-15
Clearing all FIB or AIB entries

Clear all FIB or AIB entries as part of a clean up operation, or when troubleshooting.

Step	Action
To clear all FIB entries	
1	Clear all FIB entries: `ip fib flush`
To clear all AIB entries	
2	Clear all AIB entries `ip aib flush`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 8-16
Enabling logging of FIB or AIB events

Enable logging of FIB or AIB events for troubleshooting.

Step	Action

To enable logging of FIB events

1 Enable logging of FIB events:

```
ip fib debug events on
```

To enable logging of AIB events

2 Enable logging of AIB events:

```
ip aib debug events on
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 8-17
Displaying FIB information

Display FIB information for troubleshooting purposes. You can display the following information for the FIB:

- log
- interfaces
- interface IP addresses
- static resolutions

Step	Action

To display the FIB log

1 Display the FIB log:

```
ip fib log show
```

To display the FIB interfaces

2 Display the FIB interfaces:

```
ip fib interfaces show [summary] [details]
```

where

summary	displays a summary of FIB information for interfaces.
details	displays detailed FIB information for interfaces.

To display the FIB interface IP addresses

3 Display the FIB interface IP addresses:

```
ip fib interfaceIpAddresses show [summary] | [details]
```

where

summary	displays a summary of FIB information for interface IP addresses.
details	displays detailed FIB information for interface IP addresses.

To display the FIB static resolutions

4 Display the FIB static resolutions:

```
ip fib staticResolutions show [summary] | [details]
```

where

summary	displays a summary of FIB information for static resolutions.
details	displays detailed FIB information for static resolutions.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 8-18
Adding a static ARP entry

ARP maps an IP address to a physical MAC address. Entries can be learned dynamically, such as when a device sends an ARP request or ping, but for the entries to be permanent, as in the case of using a device as a static reflector for RFC 2544, you can add a static entry.

Step	Action
1	Add a static ARP entry: `arp static add {destination <IpAddress>} {mac <MacAddress>} {ip-interface <Interface>}` where destination <IpAddress> is the destination host IP address. mac <MacAddress> is the destination MAC address. ip-interface <Interface> is the IP interface routed to the destination. **—end—**

Copyright© 2012-2013 Ciena® Corporation

Procedure 8-19
Removing static ARP entries

Remove static ARP entries when they are no longer needed. You can remove

- one static ARP entry
- all static ARP entries

Step	Action

To remove a static ARP entry

1 Remove a static ARP entry.

```
arp static remove {destination <IpAddress>} {ip-interface
<Interface>}
```

where

destination <IpAddress>	is the destination host IP address.
ip-interface <Interface>	is the IP interface routed to the destination.

To remove all static ARP entries

2 Remove all static ARP entries.

```
arp static remove-all
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 8-20
Configuring EVC ping for use case A

Configure EVC ping for a virtual switch with EPL to IP interface with virtual switch member configuration.

Step	Action
Configure NTE-1	
1	Create the VLAN: `vlan create vlan <vlan>`
2	Add a port to the VLAN: `vlan add vlan <vlan> port <port>`
3	Create an Ethernet virtual circuit: `virtual-circuit ethernet create vc <vc> vlan <vlan>`
4	Add a reserved VLAN: `virtual-switch add reserved-vlan <reserved-vlan>`
5	Create an Ethernet virtual switch: `virtual-switch ethernet create vs <vs> vc <vc>`
6	Add subscriber members to the virtual switch: `virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} [statistics <on | off>] [translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100 | 9100 | 88A8>}`
7	Create the IP interface for the virtual switch: `interface create ip-interface <ip-interface> ip <ipv4 address>/<subnet-mask> vlan <vlan>`
Configure NTE-2	
8	Create the VLAN: `vlan create vlan <vlan>`
9	Add a port to the VLAN: `vlan add vlan <vlan> port <port>`
10	Create an Ethernet virtual circuit: `virtual-circuit ethernet create vc <vc> vlan <vlan>`
11	Add a reserved VLAN: `virtual-switch add reserved-vlan <reserved-vlan>`

Copyright© 2012-2013 Ciena® Corporation

12 Create an Ethernet virtual switch:

```
virtual-switch ethernet create vs <vs> vc <vc>
```

13 Add subscriber members to the virtual switch:

```
virtual-switch ethernet add vs <vs> {port <port>} {vlan
<VLAN list>} {ip-interface <interface-name>} {encap-cos-
policy <dot1dpri-inherit | fixed | ip-prec-inherit |
phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-
dot1dpri <NUMBER: 0-7>} [statistics <on | off>]
[translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100
| 9100 | 88A8>}
```

14 Create the IP interface for the virtual switch:

```
interface create ip-interface <ip-interface> ip <ipv4
address>/<subnet-mask> vlan <vlan>
```

15 Create the IP interface:

```
interface create ip-interface <ip-interface> ip <ipv4
address>/<subnet-mask>
```

16 Add subscriber members to the virtual switch:

```
virtual-switch ethernet add vs <vs> {port <port>} {vlan
<VLAN list>} {ip-interface <interface-name>} {encap-cos-
policy <dot1dpri-inherit | fixed | ip-prec-inherit |
phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-
dot1dpri <NUMBER: 0-7>} [statistics <on | off>]
[translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100
| 9100 | 88A8>}
```

The ip-interface is operational.

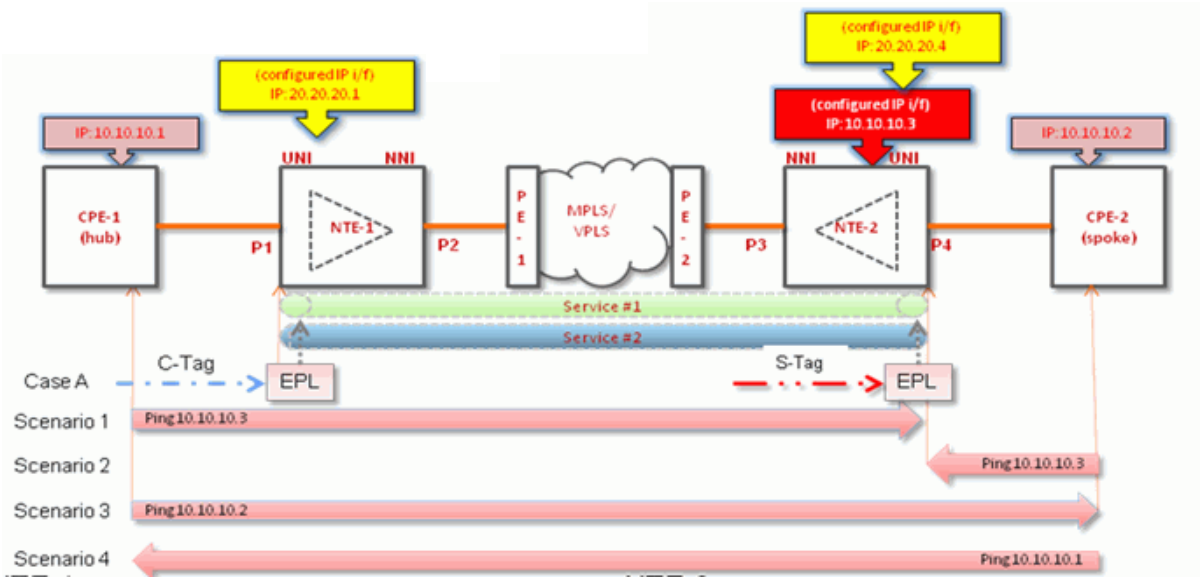
—end—

Example

Figure 8-3 shows a sample virtual switch with EPL to IP interface with virtual switch member configuration.

Copyright© 2012-2013 Ciena® Corporation

Figure 8-3
Use case A configuration example



To configure the example shown in Figure 8-3:

Configure NTE-1.

```
vlan create vlan 8
vlan add vlan 8 port 1
virtual-circuit ethernet create vc vc8 vlan 8
virtual-switch add reserved-vlan 4000
virtual-switch ethernet create vs vs8 vc vc8
virtual-switch ethernet add vs vs8 port 1
interface create ip-interface ipvs8 ip 20.20.20.1/24 vlan
8
```

Configure NTE-2.

```
vlan create vlan 8
vlan add vlan 8 port 4
virtual-circuit ethernet create vc vc8 vlan 8
virtual-switch add reserved-vlan 4000
virtual-switch ethernet create vs vs8 vc vc8
virtual-switch ethernet add vs vs8 port 4
interface create ip-interface ipvs8 ip 20.20.20.4/24 vlan
8
interface create ip-interface ip24 ip 10.10.10.3/24
virtual-switch ethernet add vs 8 ip-interface ip24 vlan 24
```

Procedure 8-21
Configuring EVC ping for use case B

Configure EVC ping for a virtual switch with EPL and untagged data virtual switch to IP interface with VLAN configuration.

Step	Action
Configure NTE-1	
1	Create the VLAN: `vlan create vlan <vlan>`
2	Add a port to the VLAN: `vlan add vlan <vlan> port <port>`
3	Create an Ethernet virtual circuit: `virtual-circuit ethernet create vc <vc> vlan <vlan>`
4	Add a reserved VLAN: `virtual-switch add reserved-vlan <reserved-vlan>`
5	Create an Ethernet virtual switch: `virtual-switch ethernet create vs <vs> vc <vc>`
6	Add subscriber members to the virtual switch: `virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} [statistics <on | off>] [translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100 | 9100 | 88A8>}`
7	Set the virtual switch for untagged data frames: `port set port <port> untagged-data-vs <vs>`
Configure NTE-2	
8	Create the VLAN: `vlan create vlan <vlan>`
9	Add a port to the VLAN: `vlan add vlan <vlan> port <port>`
10	Create an Ethernet virtual circuit: `virtual-circuit ethernet create vc <vc> vlan <vlan>`
11	Add a reserved VLAN: `virtual-switch add reserved-vlan <reserved-vlan>`

Copyright© 2012-2013 Ciena® Corporation

12 Create an Ethernet virtual switch:

```
virtual-switch ethernet create vs <vs> vc <vc>
```

13 Add subscriber members to the virtual switch:

```
virtual-switch ethernet add vs <vs> {port <port>} {vlan
<VLAN list>} {ip-interface <interface-name>} {encap-cos-
policy <dot1dpri-inherit | fixed | ip-prec-inherit |
phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-
dot1dpri <NUMBER: 0-7>} [statistics <on | off>]
[translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100
| 9100 | 88A8>}
```

14 Create the IP interface for the virtual switch and EPL:

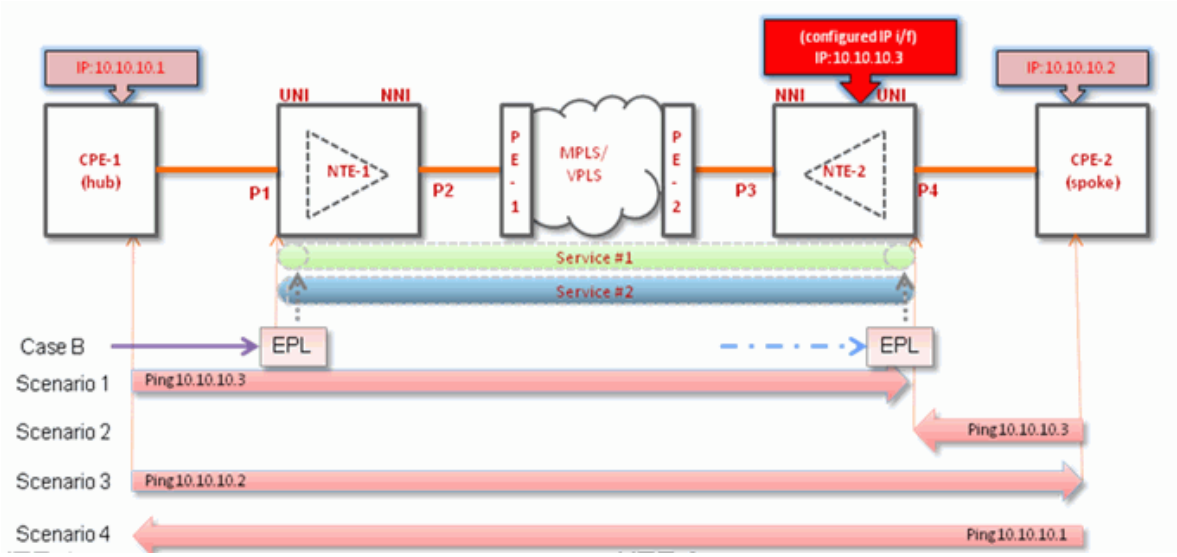
```
interface create ip-interface <ip-interface> ip <ipv4
address>/<subnet-mask> vlan <vlan>
```

—end—

Example

Figure 8-4 shows a sample virtual switch with EPL and untagged data virtual switch to IP interface with VLAN configuration.

Figure 8-4
Use case B configuration example



To configure the example shown in Figure 8-4:

Configure NTE-1.

```
vlan create vlan 8
vlan add vlan 8 port 1
virtual-circuit ethernet create vc vc8 vlan 8
```


Copyright© 2012-2013 Ciena® Corporation

```
virtual-switch add reserved-vlan 100
virtual-switch ethernet create vs vs8 vc vc8
virtual-switch ethernet add vs vs8 port 1
port set port 1 untagged-data-vs vs8
```

Configure NTE-2.

```
vlan create vlan 8
vlan add vlan 8 port 4
virtual-circuit ethernet create vc vc8 vlan 8
virtual-switch add reserved-vlan 100
virtual-switch ethernet create vs vs8 vc vc8
virtual-switch ethernet add vs vs8 port 4
interface create ip-interface ipvs8 ip 10.10.10.3/24 vlan
8
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 8-22
Configuring EVC ping for use case C

Configure EVC ping for a virtual switch with EVPL to IP interface with VLAN configuration.

Step	Action
Configure NTE-1	
1	Create the VLAN: `vlan create vlan <vlan>`
2	Add a port to the VLAN: `vlan add vlan <vlan> port <port>`
3	Create an Ethernet virtual circuit: `virtual-circuit ethernet create vc <vc> vlan <vlan>`
4	Add a reserved VLAN: `virtual-switch add reserved-vlan <reserved-vlan>`
5	Create an Ethernet virtual switch: `virtual-switch ethernet create vs <vs> vc <vc>`
6	Add subscriber members to the virtual switch (EVPL): `virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} [statistics <on | off>] [translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100 | 9100 | 88A8>}`
Configure NTE-2	
7	Create the VLAN: `vlan create vlan <vlan>`
8	Add a port to the VLAN: `vlan add vlan <vlan> port <port>`
9	Create an Ethernet virtual circuit: `virtual-circuit ethernet create vc <vc> vlan <vlan>`
10	Add a reserved VLAN: `virtual-switch add reserved-vlan <reserved-vlan>`
11	Create an Ethernet virtual switch: `virtual-switch ethernet create vs <vs> vc <vc>`

Copyright© 2012-2013 Ciena® Corporation

12 Add subscriber members to the virtual switch:

```
virtual-switch ethernet add vs <vs> {port <port>} {vlan
<VLAN list>} {ip-interface <interface-name>} {encap-cos-
policy <dot1dpri-inherit | fixed | ip-prec-inherit |
phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-
dot1dpri <NUMBER: 0-7>} [statistics <on | off>]
[translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100
| 9100 | 88A8>}
```

13 Create the IP interface for the virtual switch and EPL:

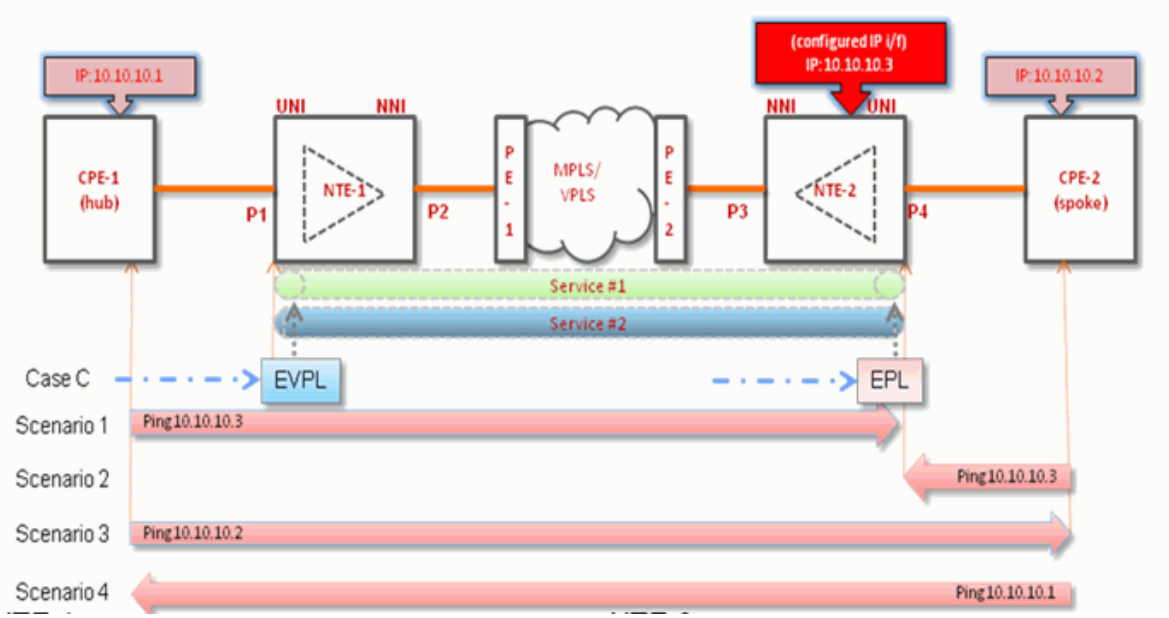
```
interface create ip-interface <ip-interface> ip <ipv4
address>/<subnet-mask> vlan <vlan>
```

—end—

Example

Figure 8-5 shows a sample virtual switch with EVPL to IP interface with VLAN configuration.

Figure 8-5
Use case C configuration example



To configure the example shown in Figure 8-5:

Configure NTE-1.

```
vlan create vlan 8
vlan add vlan 8 port 1
virtual-circuit ethernet create vc vc8 vlan 8
virtual-switch add reserved-vlan 100
```


Copyright© 2012-2013 Ciena® Corporation

```
virtual-switch ethernet create vs vs8 vc vc8
virtual-switch ethernet add vs vs8 port 1
virtual-switch ethernet add vs vs8 port 1 vlan 24
translate-tag 0
```

Configure NTE-2.

```
vlan create vlan 8
vlan add vlan 8 port 4
virtual-circuit ethernet create vc vc8 vlan 8
virtual-switch add reserved-vlan 100
virtual-switch ethernet create vs vs8 vc vc8
virtual-switch ethernet add vs vs8 port 4
interface create ip-interface ipvs8 ip 10.10.10.3/24 vlan
8
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 8-23
Configuring EVC ping for use case D

Configure EVC ping for virtual switch with EPL and untagged data virtual switch to IP interface with virtual switch member configuration.

Step	Action
Configure NTE-1	
1	Create the VLAN: `vlan create vlan <vlan>`
2	Add a port to the VLAN: `vlan add vlan <vlan> port <port>`
3	Create an Ethernet virtual circuit: `virtual-circuit ethernet create vc <vc> vlan <vlan>`
4	Add a reserved VLAN: `virtual-switch add reserved-vlan <reserved-vlan>`
5	Create an Ethernet virtual switch: `virtual-switch ethernet create vs <vs> vc <vc>`
6	Set the virtual switch for untagged data frames and set push/pop for the specified VLAN ID for untagged data frames: `port set port <port> untagged-data-vs <vs> untagged-data-vid <vlan>`
7	Create the ip-interface for the virtual switch: `interface ip create ip-interface <ip-interface> ip <ipv4 address>/<subnet-mask> vlan <vlan>`
Configure NTE-2	
8	Create the VLAN: `vlan create vlan <vlan>`
9	Add a port to the VLAN: `vlan add vlan <vlan> port <port>`
10	Create an Ethernet virtual circuit: `virtual-circuit ethernet create vc <vc> vlan <vlan>`
11	Add a reserved VLAN: `virtual-switch add reserved-vlan <reserved-vlan>`
12	Create an Ethernet virtual switch: `virtual-switch ethernet create vs <vs> vc <vc>`

Copyright© 2012-2013 Ciena® Corporation

13 Add subscriber members to the virtual switch:

```
virtual-switch ethernet add vs <vs> {port <port>} {vlan
<VLAN list>} {ip-interface <interface-name>} {encap-cos-
policy <dot1dpri-inherit | fixed | ip-prec-inherit |
phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-
dot1dpri <NUMBER: 0-7>} [statistics <on | off>]
[translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100
| 9100 | 88A8>}
```

14 Create the IP interface for the virtual switch:

```
interface create ip-interface <ip-interface> ip <ipv4
address>/<subnet-mask> vlan <vlan>
```

15 Create the IP interface for the virtual switch:

```
interface create ip-interface <ip-interface> ip <ipv4
address>/<subnet-mask> vlan <vlan>
```

16 Create the IP interface:

```
interface create ip-interface <ip-interface>
```

17 Add subscriber members to the virtual switch:

```
virtual-switch ethernet add vs <vs> {port <port>} {vlan
<VLAN list>} {ip-interface <interface-name>} {encap-cos-
policy <dot1dpri-inherit | fixed | ip-prec-inherit |
phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-
dot1dpri <NUMBER: 0-7>} [statistics <on | off>]
[translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100
| 9100 | 88A8>}
```

The ip-interface is operational.

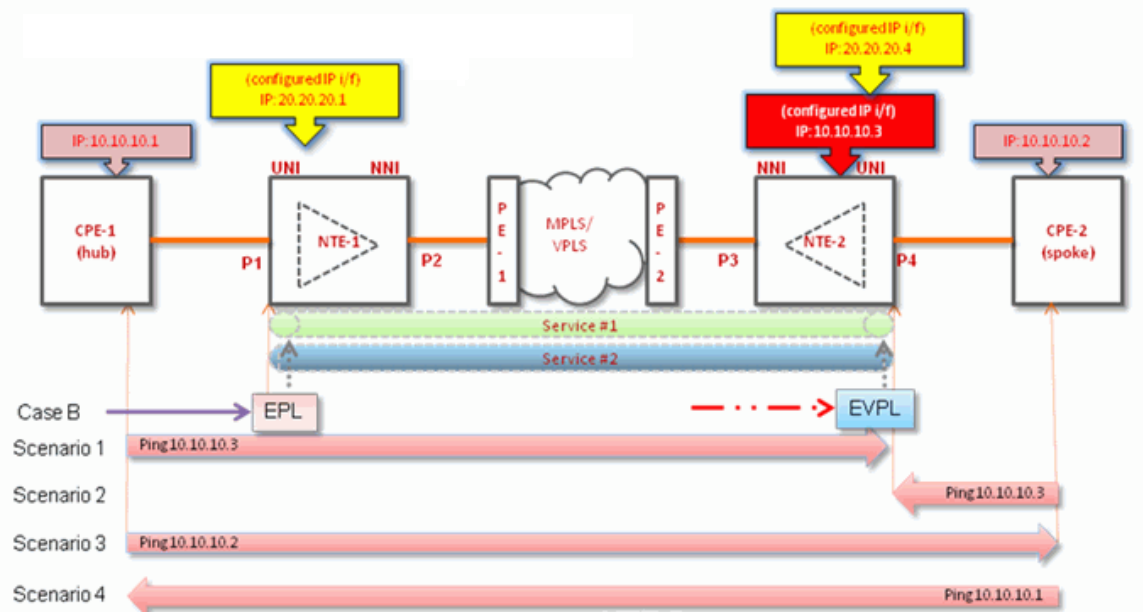
—end—

Example

Figure 8-6 shows a sample virtual switch with EPL and untagged data virtual switch to IP interface with virtual switch member configuration.

Copyright© 2012-2013 Ciena® Corporation

Figure 8-6
Use case D configuration example



To configure the example shown in Figure 8-6:

Configure NTE-1.

```
vlan create vlan 8
vlan add vlan 8 port 1
virtual-circuit ethernet create vc vc8 vlan 8
virtual-switch add reserved-vlan 100
virtual-switch ethernet create vs vs8 vc vc8
virtual-switch ethernet add vs vs8 port 1
port set port 1 untagged-data-vs vs8 untagged-data-vid 24
interface create ip-interface ipvs8 ip 20.20.20.1/24 vlan
8
```

Configure NTE-2.

```
vlan create vlan 8
vlan add vlan 8 port 4
virtual-circuit ethernet create vc vc8 vlan 8
virtual-switch add reserved-vlan 100
virtual-switch ethernet create vs vs8 vc vc8
virtual-switch ethernet add vs vs8 port 4 vlan 24
interface create ip-interface ipvs8 ip 20.20.20.4/24 vlan
8
```


Copyright© 2012-2013 Ciena® Corporation

```
interface create ip-interface ip24 ip 10.10.10.3/24
virtual-switch ethernet add vs vs8 ip-interface ip24 vlan
24
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 8-24
Configuring EVC ping for use case E

Configure EVC ping for virtual switch with EVPL to IP interface with virtual switch member configuration.

Step	Action
Configure NTE-1	
1	Create the VLAN: `vlan create vlan <vlan>`
2	Add a port to the VLAN: `vlan add vlan <vlan> port <port>`
3	Create an Ethernet virtual circuit: `virtual-circuit ethernet create vc <vc> vlan <vlan>`
4	Add a reserved VLAN: `virtual-switch add reserved-vlan <reserved-vlan>`
5	Create an Ethernet virtual switch: `virtual-switch ethernet create vs <vs> vc <vc>`
6	Add virtual switch member attributes: `virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} [statistics <on | off>] [translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100 | 9100 | 88A8>}`
7	Create the ip-interface for the virtual switch: `interface create ip-interface <ip-interface> ip <ipv4 address>/<subnet-mask> vlan <vlan>`
Configure NTE-2	
8	Create the VLAN: `vlan create vlan <vlan>`
9	Add a port to the VLAN: `vlan add vlan <vlan> port <port>`
10	Create an Ethernet virtual circuit: `virtual-circuit ethernet create vc <vc> vlan <vlan>`
11	Add a reserved VLAN: `virtual-switch add reserved-vlan <reserved-vlan>`

Copyright© 2012-2013 Ciena® Corporation

12 Create an Ethernet virtual switch:

```
virtual-switch ethernet create vs <vs> vc <vc>
```

13 Add subscriber members to the virtual switch:

```
virtual-switch ethernet add vs <vs> {port <port>} {vlan
<VLAN list>} {ip-interface <interface-name>} {encap-cos-
policy <dot1dpri-inherit | fixed | ip-prec-inherit |
phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-
dot1dpri <NUMBER: 0-7>} [statistics <on | off>]
[translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100
| 9100 | 88A8>}
```

14 Create the IP interface for the virtual switch:

```
interface create ip-interface <ip-interface> ip <ipv4
address>/<subnet-mask> vlan <vlan>
```

15 Create the IP interface:

```
interface create ip-interface <ip-interface> ip <ipv4
address>/<subnet-mask>
```

16 Add subscriber members to the virtual switch:

```
virtual-switch ethernet add vs <vs> {port <port>} {vlan
<VLAN list>} {ip-interface <interface-name>} {encap-cos-
policy <dot1dpri-inherit | fixed | ip-prec-inherit |
phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-
dot1dpri <NUMBER: 0-7>} [statistics <on | off>]
[translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100
| 9100 | 88A8>}
```

The ip-interface is operational.

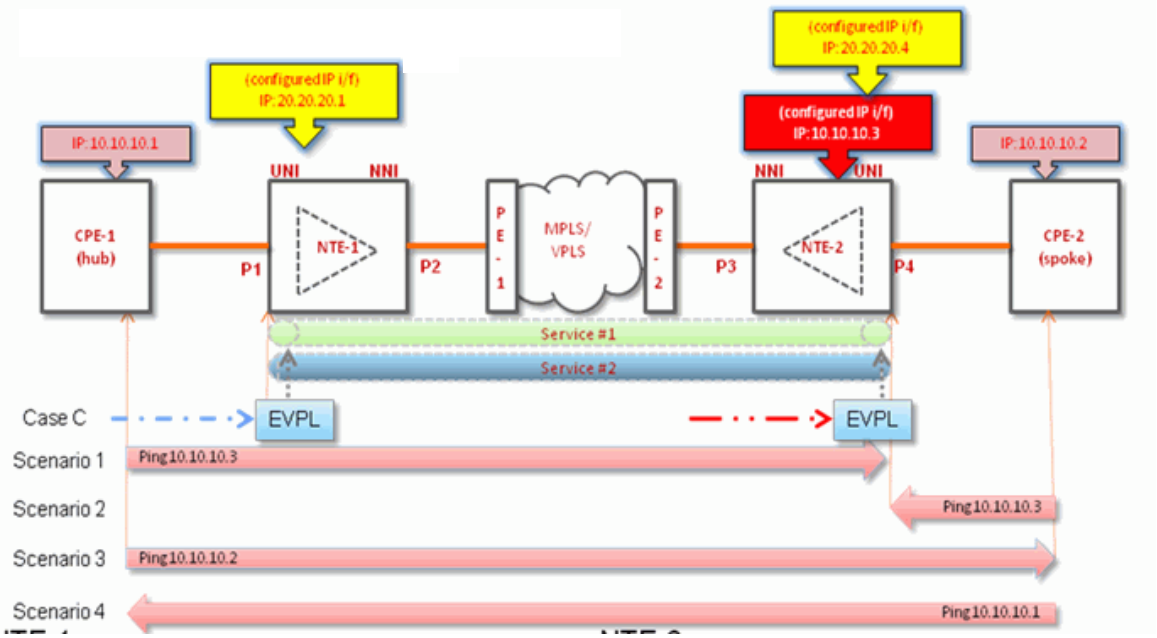
—end—

Example

Figure 8-7 shows a sample virtual switch with EVPL to IP interface with virtual switch member configuration.

Copyright© 2012-2013 Ciena® Corporation

Figure 8-7
Use case E configuration example



To configure the example shown in Figure 8-7:

Configure NTE-1.

```
vlan create vlan 8
vlan add vlan 8 port 1
virtual-circuit ethernet create vc vc8 vlan 8
virtual-switch add reserved-vlan 100
virtual-switch ethernet create vs vs8 vc vc8
virtual-switch ethernet add vs vs8 port 1 vlan 24
interface create ip-interface ipvs8 ip 20.20.20.1/24 vlan
8
```

Configure NTE-2.

```
vlan create vlan 8
vlan add vlan 8 port 4
virtual-circuit ethernet create vc vc8 vlan 8
virtual-switch add reserved-vlan 100
virtual-switch ethernet create vs vs8 vc vc8
virtual-switch ethernet add vs vs8 port 4 vlan 24
interface create ip-interface ipvs8 ip 20.20.20.4/24 vlan
8
```

```
interface create ip-interface ip24 ip 10.10.10.3/24
virtual-switch ethernet add vs vs8 ip-interface ip24 vlan
24
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 8-25
Displaying static ARP entries

Display static ARP entries.

Step	Action
1	Display static ARP entries: `arp static show`

—end—

Example

The following example shows sample output for the arp static show command.

```
> arp static show
+----------------+------------------+--------------------+
| DestinationIp  |    DMAC          | ifName (ifIndex)   |
+----------------+------------------+--------------------+
|1.1.1.2         |00:02:5a:01:b3:c6 |Ref_Test        (5 )|
+----------------+------------------+--------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 8-26
Displaying AIB information

After a static ARP entry is added, the adjacency information is stored in the AIB.

Step	Action
1	Display the AIB: `ip aib show [summary] | [details]` where summary — idisplays a summary of AIB information. details — displays detailed AIB information.

Example

The following example shows sample output for the ip aib show command.

```
> ip aib show
+---------------+-------+------------------+--------+
| NexthopIp     |ifIndex|   Dmac           | ObjIdx |
+---------------+-------+------------------+--------+
|1.1.1.2        |5      |00:02:5a:01:b3:c6 |4       |
+---------------+-------+------------------+--------+
```


Copyright© 2012-2013 Ciena® Corporation

Copyright© 2012-2013 Ciena® Corporation

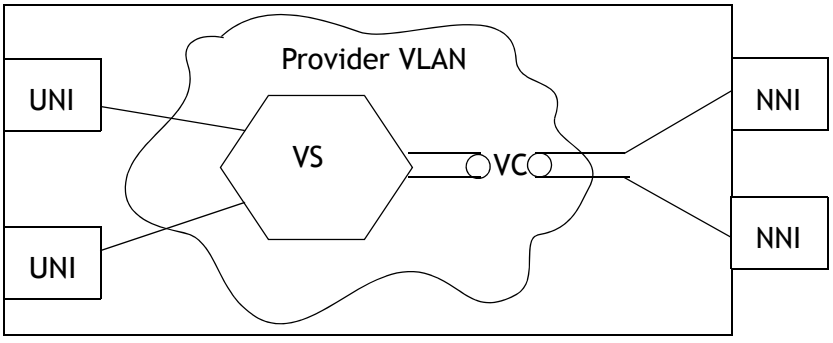
MEF L2 VPN configuration

This chapter describes Layer 2 Virtual Private Networks and how to configure them:

- "Virtual switch"
- "Private forwarding groups"
- "Virtual circuits"
- "Provider VLAN"
- "Reserved VLAN"
- "Q-in-Q encapsulation"
- "EVPL CoS"
- "EVPL Bundling"
- "VLAN translation for virtual switches"

Q-in-Q is implemented by means of the attachment of a Virtual Circuit (defining the Provider VLAN) to a Virtual Switch (defining the Subscriber VLANs), as shown in Figure 9-1. The customer perspective is that their connection to the service provider is a direct connection using a private LAN between geographically separated sites.

Figure 9-1
Provider VLAN Overview Using Q-in-Q



Copyright© 2012-2013 Ciena® Corporation

Note: In order to use the 802.1ad provider bridging feature, you need to install the Advanced Ethernet (AE) license key. License keys can be purchased by contacting Ciena customer support.

Virtual switch

A virtual switch (VS) is a Layer 2 forwarding domain between a customer bundle (multiple per-port per-VLAN entries) or per-port entries and the network transport interfaces. Ethernet switches have long been governed by the use of VLANs. Traffic tagged with a specific VLAN is forwarded to other VLAN member ports. However, in geographically separated networks, the same VLAN ID may be used by more than one subscriber group, but for security and administrative purposes, the VLAN must remain separated.

A virtual switch essentially separates the forwarding plane from the data plane, thus the virtual switch is customer-centric, while the virtual circuit is service-centric.

A virtual switch should be configured whenever there is a need for a Point-to-Point Ethernet Virtual Private LAN solution.

Note: When a port is added to a virtual switch, that port is automatically removed from its external VLANs. Therefore, any VLAN add or VLAN remove involving that port is not included in the Configuration file.

The number of VSs and members supported for Q-in-Q depends upon the platform capabilities as shown in Table 9-1

Table 9-1
Q-in-Q virtual switch capabilities

Platform	VSs	VS members per port	VS members per switch
3902, 3911, 3920	64	512	512
3916, 3930, 3931	256	1024	1024
3940, 5140	128	512	512
3960	512	2048	2048
5150	1024	3584	3584

Note 1: For 3940, 3960, 5140, and 5150 platforms, the number of VSs and VS members per port and switch supported for Q-in-Q are reduced by the number of VSs and VS members per port and switch configured for Provider Backbone Bridge Traffic Engineering (PBB-TE).

Copyright© 2012-2013 Ciena® Corporation

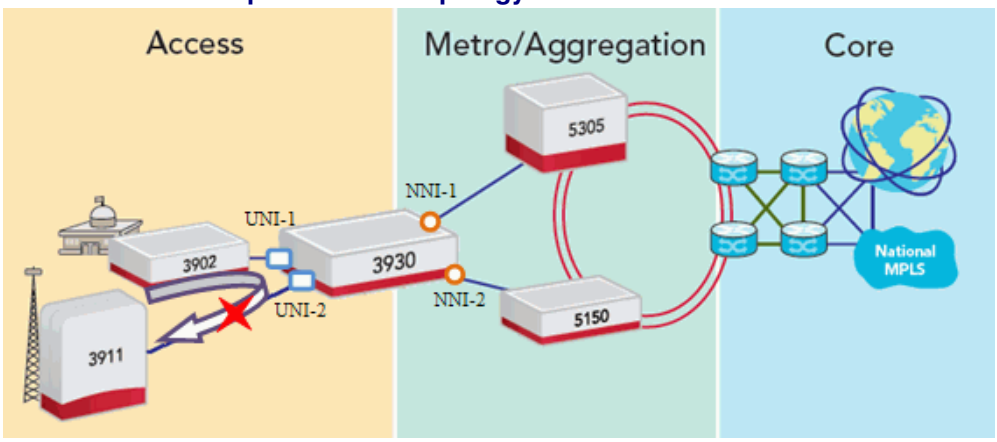
Note 2: In SAOS 4.x and 6.x, Virtual Circuits and Virtual Switches have a 1:1 relationship. A VC can be a member of one and only one VS and vice versa. This is different from SAOS 5.x, where a VS can support multiple service interfaces.

Private forwarding groups

Private forwarding groups (PFG) restrict forwarding of L2 traffic between groups of ports within a shared VLAN or virtual switch. PFGs resolve traffic loops for networks that do not use loop prevention techniques such as Spanning Tree Protocol (STP). PFGs are also used to restrict traffic forwarding between customer-facing ports within a particular forwarding domain.

Figure 9-2 shows an example where the core network connects to the customer's access network by means of two metro aggregation devices. All downlink ports on the metro devices are connected to a 3930 switch. The 3930 switch is used to provide services to two customer networks by means of a 3902 switch and a 3911 switch.

Figure 9-2
Multi UNI/NNI example: Network topology

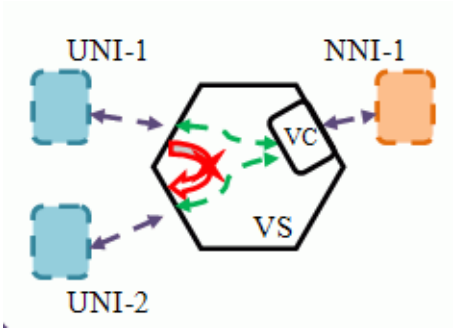


In this configuration, both customer ends are provisioned for core access, but at the same time communication is restricted between customer sites. This can be accomplished by defining a set forwarding rules for specific groups of interfaces within a forwarding domain. By separating NNIs and UNIs into two forwarding groups, or PFGs, a set of forwarding rules can be applied to each interface group. These forwarding rules can then be used to prevent traffic from being switched between customer-facing interfaces while having more than one customer provisioned under the same service instance.

Figure 9-3 illustrates the internal configuration of the 3930 switch shown in Figure 9-2.

Copyright© 2012-2013 Ciena® Corporation

Figure 9-3
Internal configuration of 3930 switch



Interfaces on 39XX/51XX platforms are divided into two PFGs: group A and group B. Each PFG is a set of ports that obeys the same forwarding rules or forwarding policy. Certain group forwarding policies are configurable and can be set such that ports within the same group do not forward traffic to each other.

Policy sets

The PFG feature uses policy sets, which house the port-based forwarding groups and rules for VLANs and VLAN-based virtual switches. Configured services on a single platform can inherit one of two global policy sets:

- management policy set, which is a collection of forwarding policies to be used exclusively on a device's current management VLAN.
- universal policy set, which is a collection of forwarding policies that all VLAN-based services can use, with the exception of the current management VLAN.

Each policy set contains two distinct private forwarding groups (PFGs), group A and group B. All interfaces on a device are divided into one of these two PFGs for both the management and universal policy sets. This means that a port can be a member of PFG A for the management policy set but be a member of PFG B for the universal policy set. PFG membership of a particular port is configured by means of the CLI and SNMP interfaces.

There is a forwarding policy, or set of rules, for every defined PFG which is applied to all ports within that particular forwarding group. A forwarding policy takes on one of the following values:

- A: can only forward to port members of private forwarding group A
- B: can only forward to port members of private forwarding group B
- A,B: can forward to port members of both private forwarding group A and B

Copyright© 2012-2013 Ciena® Corporation

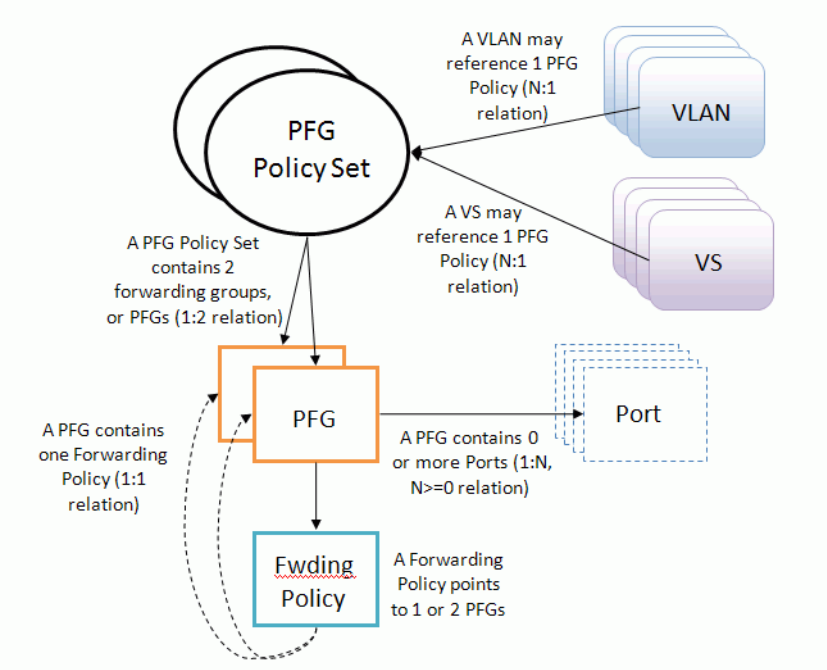
In the default configuration, all ports in PFG A can forward frames to all other ports on the switch, including ports within the same forwarding group. Ports in PFG B can only forward frames to ports in PFG A. This policy set is applied to a device's current management VLAN.

When the management VLAN changes, the new forwarding domain inherits this policy set, while the previous VLAN adopts the universal policy set. When the management VLAN is changed, the administrative status of PFG on this forwarding domain becomes disabled, that is, no forwarding rules are applied. Note that the default management VLAN is VLAN 127.

The universal policy set is applied to all VLAN-based services, that is, VLANs and VLAN-based virtual switches, excluding the management VLAN. This policy set is also configurable by means of the CLI and SNMP interfaces.

Figure 9-4 shows the structure of the policy set.

Figure 9-4
Policy set structure



Port membership

Physical ports on 39XX/51XX switches are divided into PFGs upon system boot. Each PFG policy contains its own set of forwarding groups. The PFG membership of a port within a PFG policy is configurable at any time through both the CLI and SNMP interfaces. Port membership also extends to aggregated ports through Link Aggregation Groups (LAGs).

Copyright© 2012-2013 Ciena® Corporation

When a LAG is created, the interface defaults to PFG A. Any port that is added to a LAG inherits the forwarding group of that aggregation. Similarly, when a port is removed from the aggregation, the port's original PFG membership (the forwarding group before entering into the LAG) is restored.

Table 9-2 lists default values for PFG A and PFG B. Note that the PFG feature is not supported on the 3902 switch.

Table 9-2
Platform PFG port membership default values

Platform	PFG A	PFG B
3911	1-10	--
3916	1-6	--
3920	1-12	--
3930	1-8	9,10
3931	1-8	9,10
3940	1-24	--
3960	1-10	11,12
5140	1-24	--
5150	1.1-1.36	1.37-1.48, 2.1, 2.2, 3.1, 3.2

Virtual circuits

A virtual circuit (VC) defines a secure logical connection between one or more customer endpoints. It is a service trunk, where transport services start and end (encap and decap).This association is based on an extension to IEEE 802.1ad VLAN standard referred to as double tagging (Q-in-Q). Virtual circuits encapsulate and decapsulate L2 tags over the service provider's network to allow data to be easily switched without the need for in-depth individual packet analysis or routing decisions at each network device. They can be thought of as the provider VLAN.

Copyright© 2012-2013 Ciena® Corporation

The number of VCs supported for Q-in-Q depends upon the platform capabilities as shown in Table 9-3

Table 9-3
Q-in-Q virtual circuits supported per platform

Platform	VCs
3902, 3911, 3920	64
3916, 3930, 3931	256
3940, 5140	128
3960	512
5150	1024

Note: For 3940, 3960, 5140, and 5150 platforms, the number of VCs supported for Q-in-Q are reduced by the number of VSs and VS members per port and switch configured for Provider Backbone Bridge Traffic Engineering (PBB-TE).

Provider VLAN

The provider VLAN is defined by the VLAN assigned to the virtual circuit when it is created. When a VC with a provider VLAN is attached to a VS, the provider VLAN itself becomes the double-tagged (encapsulated) VLAN. The virtual circuit references the provider VLAN, and the VC VLAN has a service port associated with it.

Note: By default, the outer Tag Protocol Identifier (TPID) of frames egressing a member port is stamped with a value of 8100 by default. This setting can be changed to 88a8 or 9100 or back to the default of 8100.

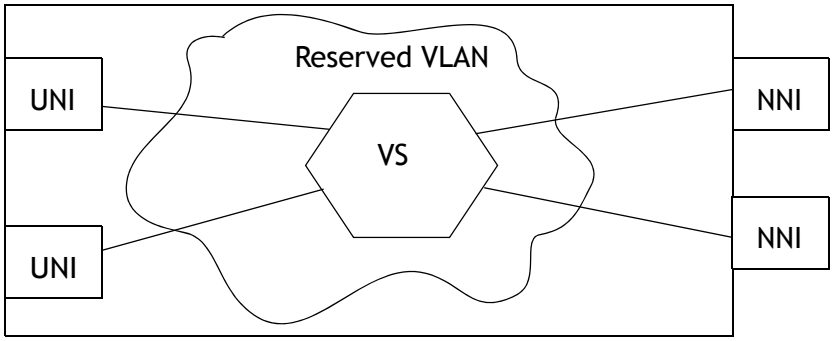
Reserved VLAN

A reserved VLAN is used internally to link a virtual switch with an internal switching domain, and acts to connect both subscriber and provider facing interfaces. A reserved VLAN must be created for each virtual switch before the virtual switch is configured. The reserved VLAN pool is part of the same 1-4094 range as the system VLAN pool. The system can automatically assign the reserved VLAN from this pool when the VS is created or you can specify a reserved VLAN when you create the VS.

Subscriber/UNI traffic can also be switched without connecting to the provider/transport network.

Copyright© 2012-2013 Ciena® Corporation

Figure 9-5
Reserved VLAN



Q-in-Q encapsulation

Encapsulating an 802.1Q VLAN tagged frame within another 802.1Q tag enables the service provider to encapsulate the customer VLAN tagging information while using VLANs within the provider network. SAOS allows Ethernet virtual circuits to be statically created. An Ethernet-VC uses an additional 802.1Q VLAN tag and includes the following:

- Service Etype (default 0x8100)
- Service VLAN ID (value 0-4095)
- Service VLAN Priority Code Point (PCP) (also, called SVLAN 802.1D priority) (value 0-7)

An Ethernet-VC has the following properties which can be configured statically:

- Name
- Provider VLAN ID and port
- Encapsulation (encap) Priority Policy (802.1D)

Q-in-Q configurations

The following configurations are supported:

- Per-port (PP) Q-in-Q: SAOS supports point-to-point Ethernet Private Line (EPL) with one EVC per UNI. Frames from port-based and port-vlan-based classified UNI ingress are forwarded to an NNI egress by adding a specified SVID, and are forwarded to a UNI egress by removing the SVID. The UNI port only accepts an untagged-data-vs configuration that matches the VS of the Per Port (EPL) member. This can be used in conjunction with the ingress-vs-filter.

Copyright© 2012-2013 Ciena® Corporation

- Per-port-per-VLAN (PPV) Q-in-Q: SAOS supports point-to-point Ethernet Virtual Private Line (EVPL) with multiple EVCs per UNI and fixed policy priority bit re-marking on ingress. Frames from port-VLAN-based classified UNI ingress are forwarded to the NNI egress by adding a specified SVID.

Frame flooding behavior

Frames received by an Ethernet virtual switch that are determined to be unknown unicast (also known as a destination lookup failure (DLF), unknown multicast, or broadcast will be flooded to the following:

- all ports that are members of the virtual circuit's VLAN, if a virtual circuit is assigned to the virtual switch
- all ports that are included as an EPL or EVPL member of the virtual switch, regardless if a customer VLAN tag exists, or if the C-VID of the frame matches the C-VID configured against the EVPL member

Q-in-Q Ethertype

The Ethertype of the outer VLAN tag can be customized. This allows compatibility with standard 802.1Q encapsulation methods like Extreme's vMAN that uses the 9100, and 802.1ad for Provider Bridging. Using the `virtual-circuit ethernet set port` command, the Ethertype value to use can be configured as described in Table 9-4.

Table 9-4
Ethertype setting (per-port)

Name	Description
8100	This indicates to use normal 802.1Q Ethertype for all frames egressing the port.
9100	This indicates to use the 0x9100 Ethertype on frames using the user configured policy described in Table 9-5.
88a8	This indicates to use the 802.1ad Ethertype on frames using the user configured policy described in Table 9-5.

Copyright© 2012-2013 Ciena® Corporation

The `virtual-circuit ethernet set port` command includes a setting that allows the operator to specify the policy or rules for applying that configured Ethertype on a per-port basis as described in Table 9-5.

Table 9-5
Ethertype policy (per-port)

Name	Description
all	When this policy is set, only Q-in-Q encapsulated frames going out the port use the configured Ethertype (see Table 9-4).
encap-only	When this policy is set, only encapsulated frames going out a port use the configured Ethertype (see Table 9-4). This affects both single and double tagged frames that have had a provider VLAN tag added to them.
vlan-tpid	When this policy is set, frames going out a port use the based on the egress-tpid value assigned to the provider VLAN.

EVPL CoS

Virtual Switch CoS Policies

Each VS has default fixed encapsulation CoS policy (encap-cos-policy) to mark the 802.1D priority value in the SVLAN tag that is pushed onto the customer frame on egress from the UNI with the configured encapsulation 802.1D priority value (encap-fixed-dot1dpri). The CoS policy for the CVLAN tag (subscriber-dot1dpri-policy) can only be set to “leave”, which does not modify the frame. Also, for decapsulation, VSs only support a “leave” CoS policy. When a VS is created the default encap-fixed-dot1dpri value is 2 and the subscriber-dot1dpri-policy is leave.

VS Member CoS Policy Override

By default, VS members inherit the CoS policy from the VS. You can configure the encap-cos-policy of the VS member to one of three settings:

- fixed — The SVLAN 802.1D value is taken from the encap-fixed-dot1dpri parameter of the VS member.
- port-inherit — The SVLAN 802.1D value is based on the resolved Cos policy of the port, which can be one of the following:
  - — fixed CoS (fixed-cos) — The SVLAN 802.1D value comes from the configured fixed resolved CoS (fixed-rcos) of the port.
  - — .1D mapped (dot1d-tag1-cos) — Value is mapped from the CVLAN 802.1D value as shown in Table 9-6.

Copyright© 2012-2013 Ciena® Corporation

Table 9-6
Default .1D Mapped CoS Mapping Table

Customer Frame 802.1D	SVLAN 802.1D
0	0
1	1
2	2
3	3
4	4
5	5
6	6
7	7

— DSCP mapped (l3-dscp-cos) - Value is mapped from the CVLAN DSCP value as shown in Table 9-7.

Table 9-7
Default DSCP Mapping Table

Customer Frame DSCP	SVLAN 802.1D
0 - 7	0
8-15	1
16-23	2
24-31	3
32-39	4
40-47	5
48-55	6
56-63	7

— vs-inherit — Indicates that the SVLAN 802.1D value is defined by the VS encap CoS policy configuration. This setting is the default.

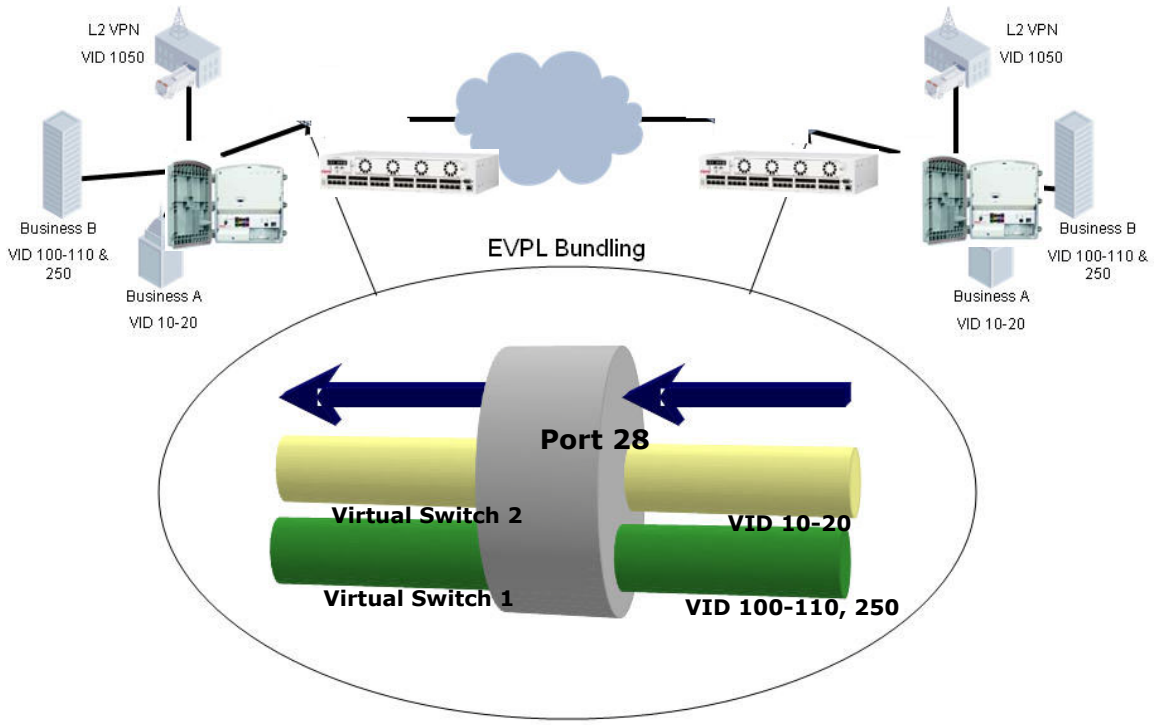
EVPL Bundling

SAOS supports EVPL bundling. EVPL bundling allows multiple customer VLAN IDs (C-VIDs) to be mapped to one or more virtual switches. It should be noted that EVPL bundling only functions where traffic is being encapsulated.

As an example, Figure 9-6 shows multiple C-VIDS ingressing on port 28. VLANs 10-20 have been mapped to virtual switch 2, and VLANs 100-110 and 250 are mapped to Ethernet-VS 1. This allows multiple C-VIDs to be mapped to two separate virtual switches.

Copyright© 2012-2013 Ciena® Corporation

Figure 9-6
Multiple C-VIDS to Multiple VS



Example

```
> virtual-switch ethernet add vs vs1 port 28 vlan 100-110,250
> virtual-switch ethernet add vs vs2 port 28 vlan 10-20
```

Note: If the port vlan-ingress-filter parameter is set to on, only frames that match a configured VLAN tag are forwarded.

VLAN translation for virtual switches

The 3916, 3930, 3931, 3960 and 5150 platforms support configurable VLAN translation to classify Q-in-Q traffic from one L2 switching domain and remap it to another.

To configure VLAN translation for Q-in-Q traffic, set the virtual switch L2 transform actions on the UNI ports and a translation VLAN tag for the virtual switch. L2 transform actions include:

Ingress push; egress pop (i-push,e-pop)

Upon ingress, the device looks up the frame's VID (which becomes the C-VID) and classifies to a virtual switch, then pushes the SVID tag onto the frame. At egress, the device looks up the frame's SVID and pops the SVID tag. This is the default L2 transform action and is supported by all platforms running this system software.

Copyright© 2012-2013 Ciena® Corporation

When a UNI port is configured with this L2 transform mode, the following rules apply to the associated virtual switch:

- EPL members can be added.
- Single or multiple EVPL members can be added.
- Multiple UNI port members can be added.
- Can be specified as the virtual switch for untagged data and untagged control frames for the port.
- L2 control frame tunneling can be enabled.
- All virtual switch members must have a translate tag of zero.

Ingress push; egress pop and stamp (i-push,e-pop:stamp)

Upon ingress, the device looks up the frame's VID (which becomes the CVID) and classifies it to a virtual switch, and pushes the SVID tag onto the frame. At egress, the device looks up the frame's SVID, pops the SVID tag, and stamps the VID with the ingress classifier CVID.

When a UNI port is configured with this L2 transform mode, the following rules apply to the associated virtual switch:

- EPL members cannot be added.
- Only one EVPL member can be added.
- Cannot add multiple UNI port members.
- Cannot be specified as the virtual switch for untagged data and untagged control frames for the port.
- L2 control frame tunneling cannot be enabled.
- All virtual switch members must have a translate tag of zero.

Note: Only the 3916, 3930, 3931, 3960 and 5150 support this L2 transform mode.

Ingress stamp; egress match, pop, and stamp (i-stamp:push,e-match-pop:stamp)

There are two scenarios supported under this mode depending on translate-tag configuration, and both scenarios can be configured under the same virtual switch at the same time.

When the translate-tag is non-zero:

- upon UNI ingress, the device classifies a cVid and stamps <cVid [+ cPri]>, and then pushes an sTag (Q-in-Q).

Copyright© 2012-2013 Ciena® Corporation

- upon UNI egress, the device classifies <sVid + cVid> for exact match, and then pops the sTag and stamps the cVid back to the original cVid used for ingress-classification.

When the translate-tag is zero:

- upon UNI ingress, the device classifies a cVid and stamps it with a new sTag <sVid [+sPri]>.
- upon UNI Egress, the device classifies an sVid for exact match, and then stamps it back to the original cVid used for ingress-classification.

At ingress, the device looks up the incoming frame's outer vid (cVid) and classifies it to a virtual switch. In addition, the device stamps the cVid with a new unique cVid value, and then pushes the sTag onto the frame if the translate-tag is not zero.

At egress, if the translate-tag is not zero, the device performs a qualified lookup on the frame's <sVid + cVid>. On match, the device pops the sTag and stamps the cVid with the original cVid specified as the ingress classifier. If the translate-tag is zero, the classification is on sTag and stamped with the original cVid back.

Due to the qualification of the <sVid + cVid> at egress and the relevant uniqueness of the translated cVid to a port, using this transform-mode constrains the configuration as follows:

- Virtual switch can contain only one UNI port specification.
- Virtual switch UNI member must contain one or more cVid specifications.
  - — Vid bundle is OK since it exists on a single port.
  - — The translate-tag value must be unique among cVids of the bundle.
  - — The translate-tag value may be used by one or more bundles on different virtual switches.
  - — Only one cVid can be configured with translate-tag zero.
- EPL is not allowed as a virtual switch member in this transform mode.

  Note: Only the 3916, 3930, 3931, 3960 and 5150 support this L2 transform mode.

Under the following conditions, traffic can unexpectedly obtain the same S- and C-VID and as a result, merge together toward the NNI direction:

- Both zero translate-tag and non-zero translate-tag members are provisioned under the same virtual switch.
- Double-tagged traffic classifies to the zero translate-tag member from the UNI side.

Copyright© 2012-2013 Ciena® Corporation

- A non-zero translate-tag member has the same value as the inner tag of the zero translate-tag traffic.

Ciena recommends that the network operator provision the non-zero translate-tag member under another virtual switch.

Copyright© 2012-2013 Ciena® Corporation

Procedure 9-1
Creating an EPL provider bridge

Create an EPL provider bridge.

Step	Action
1	Create a provider VLAN (SVID). `vlan create vlan <VlanIdList> [name <String[31]>]`
2	Add the NNI port to the provider VLAN. `vlan add vlan <VlanList> port <PortNameList>`
3	Add reserved VLANs for use when creating virtual switches. `virtual-switch add reserved-vlan <VLAN IDList>`
4	Create an Ethernet virtual circuit using the provider VLAN. `virtual-circuit ethernet create {vc <VirtualCircuitEthernetName[15]>} {vlan <Vlan>}`
5	Create the virtual switch and associate it with the virtual circuit created in step 4. `virtual-switch ethernet create vs <VirtualSwitchName> vc <VirtualCircuitEthernetName[15]> encap-fixed-dot1dpri <NUMBER: 0-7>`
6	Add UNI PORT+CVID members to the Ethernet virtual switches. `virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} [statistics <on | off>] [translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100 | 9100 | 88A8>}`

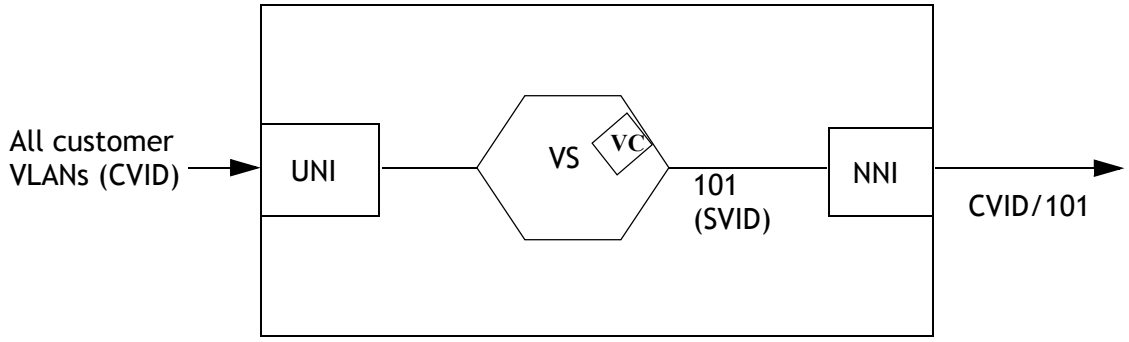
—end—

Example

Figure 9-7 shows an example of an EPL provider bridge.

Copyright© 2012-2013 Ciena® Corporation

Figure 9-7
Ethernet Private Line



The following example configures the EPL provider bridge.

```
> vlan create vlan 101
> vlan add vlan 101 port 10
> virtual-switch add reserved-vlan 1001
> virtual-circuit ethernet create vc EVC101 vlan 101
> virtual-switch ethernet create vs EVS100 vc EVC101 encap-fixed-dot1dpri 3
> virtual-switch ethernet add vs EVS100 port 1
```

Procedure 9-2
Creating an EVPL provider bridge

Create an EVPL provider bridge.

Step	Action
1	Create a provider VLAN (SVID) `vlan create vlan <VlanIdList> [name <String[31]>]`
2	Add the NNI port to the VLAN. `vlan add vlan <VlanList> port <PortNameList>`
3	Create an Ethernet virtual circuit using the provider VLAN. `virtual-circuit ethernet create {vc <VirtualCircuitEthernetName[15]>} {vlan <Vlan>}`
4	Add reserved VLANs for use when creating virtual switches. `virtual-switch add reserved-vlan <VLAN IDList>`
5	Create Ethernet virtual switches using the Ethernet virtual circuit. `virtual-switch ethernet create vs <VirtualSwitchName> vc <VirtualCircuitName>`
6	Add UNI PORT+CVID members to the Ethernet virtual switches. `virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} [statistics <on | off>] [translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100 | 9100 | 88A8>}`

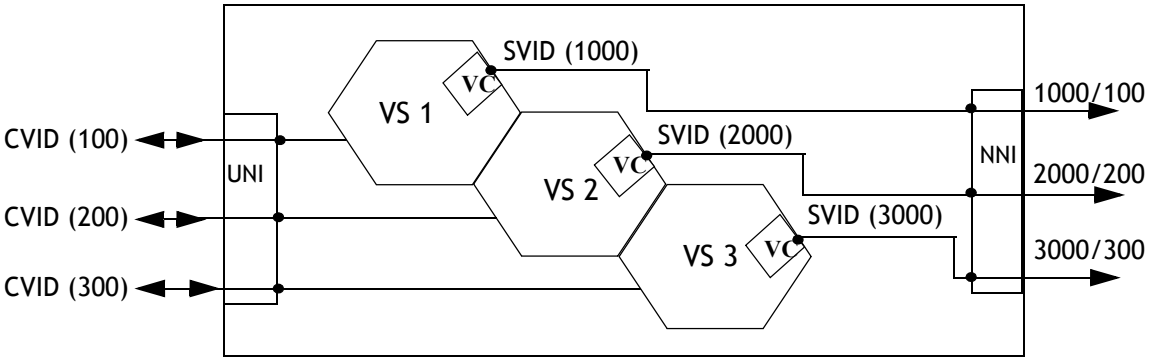
—end—

Example

Figure 9-8 shows an example of an EVPL provider bridge.

Copyright© 2012-2013 Ciena® Corporation

Figure 9-8
EVPL Provider Bridging



The following example creates the configuration.

Create a provider VLAN (SVID) and add the NNI port to the VLAN:

```
> vlan create vlan 1000
> vlan create vlan 2000
> vlan create vlan 3000
> vlan add vlan 1000,2000,3000 port 8
```

Create an Ethernet virtual circuit using the provider VLAN.

```
> virtual-circuit ethernet create vc vc1 vlan 1000
> virtual-circuit ethernet create vc vc2 vlan 2000
> virtual-circuit ethernet create vc vc3 vlan 3000
```

Add reserved VLANs for use when creating virtual switches.

```
> virtual-switch add reserved-vlan 4000,4001,4002
```

Create Ethernet virtual switches using the Ethernet virtual circuit.

```
> virtual-switch ethernet create vs vs1 vc vc1
> virtual-switch ethernet create vs vs2 vc vc2
> virtual-switch ethernet create vs vs3 vc vc3
```

Add UNI PORT+CVID members to the Ethernet virtual switches.

```
> virtual-switch ethernet add vs vs1 port 1 vlan 100
> virtual-switch ethernet add vs vs2 port 1 vlan 200
> virtual-switch ethernet add vs vs3 port 1 vlan 300
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 9-3
Configuring the fixed encapsulation priority value

Configure the fixed encapsulation priority value. You can configure the fixed encapsulation priority value for

- a virtual switch upon creation
- an existing virtual switch

Step	Action

To configure the fixed encapsulation priority value for a virtual switch upon creation

1 Configure the fixed encapsulation priority value for a virtual switch upon creation:

```
virtual-switch ethernet create vs <String> encap-fixed-
dot1dpri <NUMBER:0-7>
```

To configure the fixed encapsulation priority value for an existing virtual switch

2 Configure the fixed encapsulation priority value for an existing virtual switch:

```
virtual-switch ethernet set vs <vs> {decap-cos-policy
<fixed | leave | tunnel-inherit | vc-inherit>} {decap-
fixed-dot1dpri <NUMBER: 0-7>} {description <String[128]>}
{encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-
inherit | phbg-inherit>} {encap-fixed-dot1dpri <NUMBER:
0-7>} {subscriber-dot1dpri-policy <leave | provider-
inherit>} {vc <Virtual Circuit Name>} [ip-iterface
<interface-name>]
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 9-4
Setting the CoS policy when adding VS members

Set the CoS policy when adding VS members.

Step	Action
1	Set the CoS policy when adding VS members: `virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} [statistics <on | off>] [translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100 | 9100 | 88A8>}`
2	Set the CoS policy for existing VS members: `virtual-switch ethernet set vs <vs> {decap-cos-policy <fixed | leave | tunnel-inherit | vc-inherit>} {decap-fixed-dot1dpri <NUMBER: 0-7>} {description <String[128]>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} {subscriber-dot1dpri-policy <leave | provider-inherit>} {vc <Virtual Circuit Name>} [ip-iterface <interface-name>]`

—end—

Example

The following example sets the CoS policy when adding VS members.

```
virtual-switch ethernet add vs vs150 port 5 vlan 151
encap-cos-policy fixed encap-fixed-dot1dpri 6
```

The following example sets the CoS policy for existing VS members.

```
virtual-switch ethernet set port 5 vs vs150 encap-cos-
policy vs-inherit
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 9-5
Handling ingress untagged frames

When a UNI port is set to accept all frame types, untagged data frames and untagged control frames are allowed to ingress. If desired, you can configure the port to forward these frames to specific virtual switches with the untagged-ctrl-vs or untagged-data-vs attributes. Also, you can configure the port's untagged-data-vid attribute to add (push) a CVID tag on untagged data frames on ingress. In addition to the CVID, the tag can include the port's ingress fixed .1d priority. These frames are encapsulated with the CVID tag and a header that includes the Service VLAN ID (SVID). After processing, the frames are decapsulated, and the CVID tag is retained on egress.

Note 1: The untagged-data-vs and untagged-ctrl-vs attributes do not apply to MPLS virtual switches. All untagged traffic is directed to the MPLS EPL virtual switch on that port.

Note 2: On the 3902, 3911, 3920, 3940 and 5140 platforms, the untagged-data-vid attribute pushes the CVID on ingress but does not pop the CVID on egress.

Step	Action
1	Create a provider VLAN (SVID) `vlan create vlan <VlanIdList> [name <String[31]>]`
2	Add reserved VLANs for use when creating virtual switches. `virtual-switch add reserved-vlan <VLAN IDList>`
3	Create an Ethernet virtual circuit using the provider VLAN. `virtual-circuit ethernet create {vc <VirtualCircuitEthernetName[15]>} {vlan <Vlan>}`
4	Create the virtual switch and associate it with the virtual circuit: `virtual-switch ethernet create vs <VirtualSwitchName[15]> vc <VirtualCircuitName[15]>`
5	Add the UNI port to the virtual switch. `virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} [statistics <on | off>] [translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100 | 9100 | 88A8>}`

Copyright© 2012-2013 Ciena® Corporation

6 Set UNI port attributes:

```
port set port <PortNameList> acceptable-frame-type
<all|taggedonly> untagged-data-vs <VirtualSwitchName>
untagged-data-vid <VlanId> fixed-rcos <NUMBER: 0-7]
```

where

port <PortNameList>	is the port list
acceptable-frame-type <all|taggedonly>	is the acceptable frame types mode
untagged-data-vs <VirtualSwitchName>	is the virtual switch for untagged data frames
untagged-data-vid <VlanId>	is the VLAN ID for untagged data frames ***Note:*** On the 3902, 3911, 3920, 3940 and 5140 platforms, the untagged-data-vid attribute pushes the VLAN ID on ingress but does not pop the VLAN ID on egress.
fixed-rcos <NUMBER: 0-7]	is the Fixed Resolved CoS value

—end—

Example

The following example sets the following UNI port attributes:

- accept all frame types
- allow untagged data frames and untagged control frames to ingress
- add (push) a CVID tag on untagged data frames on ingress and include the port's ingress fixed .1d priority

```
> vlan create vlan 555
> virtual-switch add reserved-vlan 1001
> virtual-circuit ethernet create vc vc555 vlan 555
> virtual-switch ethernet create vs vsX vc vc555
> virtual-switch ethernet add vs vsX port 1
> port set port 1 acceptable-frame-type all untagged-data-vs vsX
untagged-data-vid 555 fixed-rcos 5
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 9-6
Setting the L2 transform action on a port

Set the L2 transform action on a port.

To enable transform actions other than the default on the 3916, 3930, 3931, 3960 and 5150 platforms, assign resources for the feature.

Note: For more information regarding resource management, see "Hardware resource management" on page 4-1.

Step	Action

1 Set the L2 transform action on a port:

```
port set port <PortName> vs-l2-transform <i-push,e-pop|i-push,e-pop:stamp|i-stamp:push,e-match-pop:stamp>
```

To assign resources for the feature:

2 Set the pool count for classifier resources to 512.

```
resource-manager pool set resource classifier feature vs-enhanced-l2-transform count 512
```

3 Set the pool count for meter resources to 1024.

```
resource-manager pool set resource meter feature vs-enhanced-l2-transform count 1024
```

4 Set the pool count for counter resources to 512.

```
resource-manager pool set resource counter feature vs-enhanced-l2-transform count 512
```

5 Validate the configuration.

```
resource-manager validate
```

6 Save the configuration.

```
configuration save
```

7 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 9-7
Creating an i-push, e-pop Q-in-Q VS configuration

Create an i-push, e-pop Q-in-Q configuration.

Step	Action
1	Set the virtual switch layer 2 transform action to support VLAN translation for virtual switches on the first port: `port set port <PortNameList> vs-l2-transform i-push,e-pop` where <PortNameList> is the name of the first port
2	Set the virtual switch layer 2 transform action to support VLAN translation for virtual switches on the second port: `port set port <PortNameList> vs-l2-transform i-push,e-pop` where <PortNameList> is the name of the second port
3	Create a VLAN: `vlan create vlan <VlanIdList> [name <String[31]>]`
4	Add ports to the VLAN created in step 3. `vlan add vlan <VlanList> port <PortNameList>`
5	Reserve VLANs for use in creating Virtual Switch Instances: `virtual-switch add reserved-vlan <VLAN IDList>`
6	Create a static Ethernet virtual circuit for the specified VLAN object: `virtual-circuit ethernet create {vc <VirtualCircuitEthernetName[15]>} {vlan <Vlan>}`
7	Create a virtual switch instance and associate a virtual circuit: `virtual-switch ethernet create vs <VirtualSwitchName[15]> [vc <VirtualCircuitName[15]>] reserved-vlan <VLAN IDList>`
8	Add an attachment circuit: `virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} [statistics <on | off>] [translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100 | 9100 | 88A8>}`

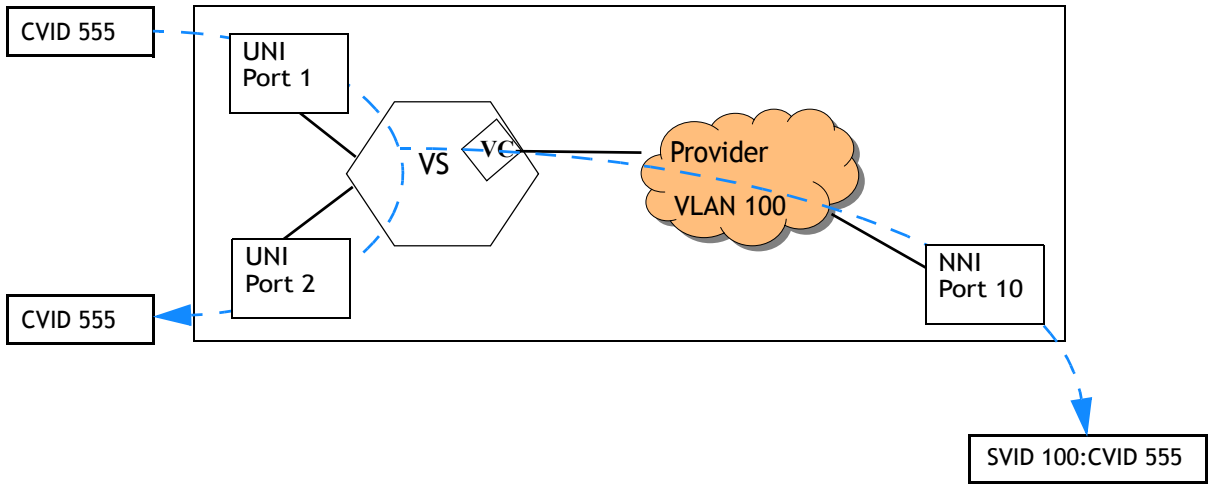
—end—

Copyright© 2012-2013 Ciena® Corporation

Example

The example illustrated in Figure 9-9 shows a Q-in-Q VS configuration where traffic is classified to CVID 555 and the SVID 100 is pushed on the frame at ingress. From UNI-1 to UNI-2, the SVID is popped on egress. From UNI to NNI, frames egress with SVID 100.

Figure 9-9
VS L2 transform: i-push,e-pop for Q-in-Q



The following example creates the configuration.

```
> port set port 1 vs-l2-transform i-push,e-pop
> port set port 2 vs-l2-transform i-push,e-pop
> vlan create vlan 100
> vlan add vlan 100 port 10
> virtual-switch add reserved-vlan 4000
> virtual-circuit ethernet create vc vc100 vlan 100
> virtual-switch ethernet create vs vs100 vc vc100 reserved-vlan 4000
> virtual-switch ethernet add vs vs100 port 1 vlan 555
```

Procedure 9-8
Creating a VS configuration with UNI only with bundled CVIDs

Create a VS configuration with UNI only with bundled CVIDs.

Step	Action
1	Set the virtual switch layer 2 transform action to support VLAN translation for virtual switches on the first port: `port set port <PortNameList> vs-l2-transform i-push,e-pop` where <PortNameList> is the name of the first port
2	Set the virtual switch layer 2 transform action to support VLAN translation for virtual switches on the second port: `port set port <PortNameList> vs-l2-transform i-push,e-pop` where <PortNameList> is the name of the second port
3	Reserve VLANs for use in creating Virtual Switch Instances: `virtual-switch add reserved-vlan <VLAN IDList>`
4	Create a virtual switch instance: `virtual-switch ethernet create vs <VirtualSwitchName[15]> reserved-vlan <VLAN IDList>`
5	Add attachment circuits for the first port: `virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} [statistics <on | off>] [translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100 | 9100 | 88A8>}`
6	Add attachment circuits for the second port: `virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} [statistics <on | off>] [translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100 | 9100 | 88A8>}`

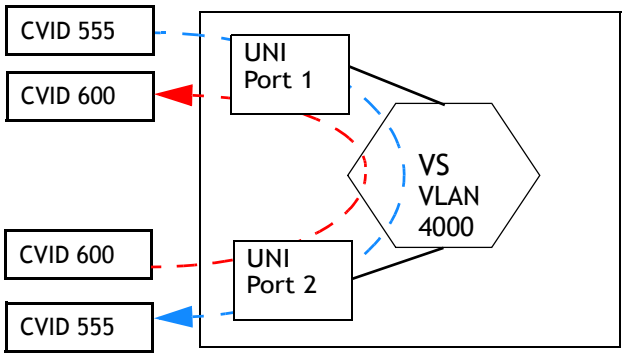
—end—

Copyright© 2012-2013 Ciena® Corporation

Example

The example illustrated in Figure 9-9 shows a VS configuration where traffic is classified to CVID 555 and 600. The SVID 4000 is pushed on the frame at ingress and popped at egress.

Figure 9-10
VS L2 transform: i-push,e-pop for UNI only with bundled CVIDs



The following example creates the configuration.

```
> port set port 1 vs-l2-transform i-push,e-pop
> port set port 2 vs-l2-transform i-push,e-pop
> virtual-switch add reserved-vlan 4000
> virtual-switch ethernet create vs vs100 reserved-vlan 4000
> virtual-switch ethernet add vs vs100 port 1 vlan 555
> virtual-switch ethernet add vs vs100 port 1 vlan 600
> virtual-switch ethernet add vs vs100 port 2 vlan 555
> virtual-switch ethernet add vs vs100 port 2 vlan 600
```

Procedure 9-9
Creating an i-push, e-pop: stamp configuration

Create an i-push, e-pop: stamp configuration.

Step	Action

To configure Device A

1 Set the virtual switch layer 2 transform action to support VLAN translation for virtual switches:

```
port set port <PortNameList> vs-l2-transform i-push,e-pop:stamp
```

2 Create a VLAN:

```
vlan create vlan <VlanIdList> [name <String[31]>]
```

3 Add ports to the VLAN created in step 2.

```
vlan add vlan <VlanList> port <PortNameList>
```

4 Create a static Ethernet virtual circuit for the specified VLAN object:

```
virtual-circuit ethernet create {vc <VirtualCircuitEthernetName[15]>} {vlan <Vlan>}
```

5 Reserve VLANs for use in creating Virtual Switch Instances:

```
virtual-switch add reserved-vlan <VLAN IDList>
```

6 Create a virtual switch instance and associate a virtual circuit:

```
virtual-switch ethernet create vs <VirtualSwitchName[15]> [vc <VirtualCircuitName[15]>] reserved-vlan <VLAN IDList>
```

7 Add an attachment circuit:

```
virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} [statistics <on | off>] [translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100 | 9100 | 88A8>}
```

To configure Device B

8 Set the virtual switch layer 2 transform action to support VLAN translation for virtual switches:

```
port set port <PortNameList> vs-l2-transform i-push,e-pop:stamp
```

9 Create a VLAN:

```
vlan create vlan <VlanIdList> [name <String[31]>]
```

10 Add ports to the VLAN created in step 9.

Copyright© 2012-2013 Ciena® Corporation

```
vlan add vlan <VlanList> port <PortNameList>
```

11 Create a static Ethernet virtual circuit for the specified VLAN object:

```
virtual-circuit ethernet create {vc
<VirtualCircuitEthernetName[15]>} {vlan <Vlan>}
```

12 Reserve VLANs for use in creating Virtual Switch Instances:

```
virtual-switch add reserved-vlan <VLAN IDList>
```

13 Create a virtual switch instance and associate a virtual circuit:

```
virtual-switch ethernet create vs <VirtualSwitchName[15]>
[vc <VirtualCircuitName[15]>] reserved-vlan <VLAN IDList>
```

14 Add an attachment circuit:

```
virtual-switch ethernet add vs <vs> {port <port>} {vlan
<VLAN list>} {ip-interface <interface-name>} {encap-cos-
policy <dot1dpri-inherit | fixed | ip-prec-inherit |
phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-
dot1dpri <NUMBER: 0-7>} [statistics <on | off>]
[translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100
| 9100 | 88A8>}
```

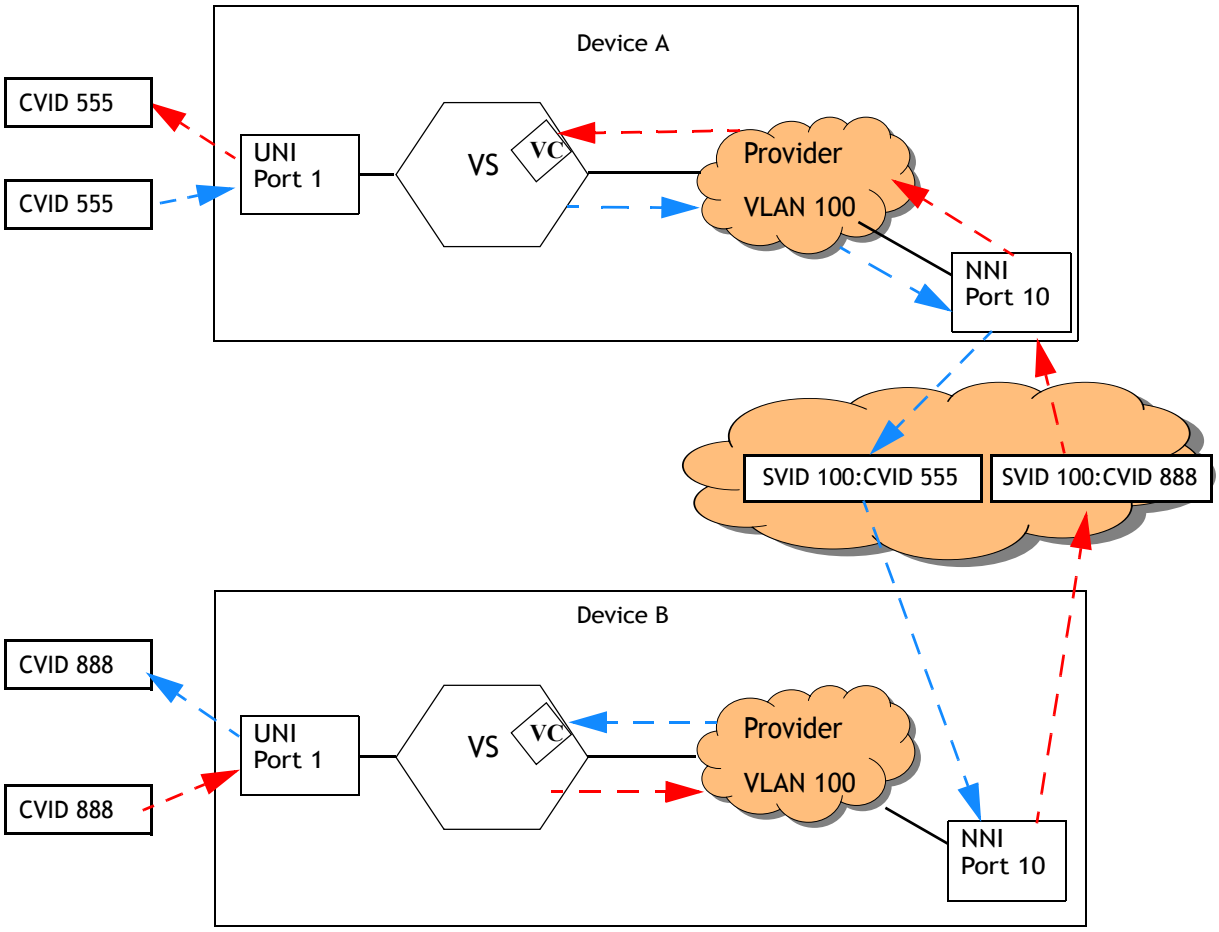
—end—

Example

In this example, ingress frames that classify to CVID 555 have the SVID 100 pushed onto them. On egress, any frames that classify to SVID 100 have the SVID popped off and the CVID 555 stamped. This configuration requires the allocation of resources as described in "Setting the L2 transform action on a port".

Copyright© 2012-2013 Ciena® Corporation

Figure 9-11
VS L2 transform: i-push,e-pop:stamp for Q-in-Q



Configure Device A.

```
port set port 1 vs-l2-transform i-push,e-pop:stamp
vlan create vlan 100
vlan add vlan 100 port 10
virtual-circuit ethernet create vc vc1 vlan 100
virtual-switch add reserved-vlan 4000
virtual-switch ethernet create vs vs100 vc vc1 reserved-vlan 4000
virtual-switch ethernet add vs vs100 port 1 vlan 555
```

Configure Device B.

```
port set port 1 vs-l2-transform i-push,e-pop:stamp
vlan create vlan 100
vlan add vlan 100 port 10
virtual-circuit ethernet create vc vc1 vlan 100
virtual-switch add reserved-vlan 4000
virtual-switch ethernet create vs vs100 vc vc1 reserved-vlan 4000
virtual-switch ethernet add vs vs100 port 1 vlan 888
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 9-10
Creating an i-stamp:push,e-match-pop:stamp configuration

Create an i-stamp:push,e-match-pop:stamp configuration.

Step	Action
1	Sets the virtual switch layer 2 transform action to support VLAN translation for virtual switches. `port set port <PortNameList> vs-l2-transform i-stamp:push,e-match-pop:stamp`
2	Create a VLAN: `vlan create vlan <VlanIdList> [name <String[31]>]`
3	Add ports to the VLAN created in step 2. `vlan add vlan <VlanList> port <PortNameList>`
4	Create a static Ethernet virtual circuit for the specified VLAN object: `virtual-circuit ethernet create {vc <VirtualCircuitEthernetName[15]>} {vlan <Vlan>}`
5	Reserve VLANs for use in creating Virtual Switch Instances: `virtual-switch add reserved-vlan <VLAN IDList>`
6	Create a virtual switch instance and associate a virtual circuit: `virtual-switch ethernet create vs <VirtualSwitchName[15]> [vc <VirtualCircuitName[15]>] reserved-vlan <VLAN IDList>`
7	Add attachment circuits: `virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} [statistics <on | off>] [translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100 | 9100 | 88A8>}`

—end—

Example

The following examples show

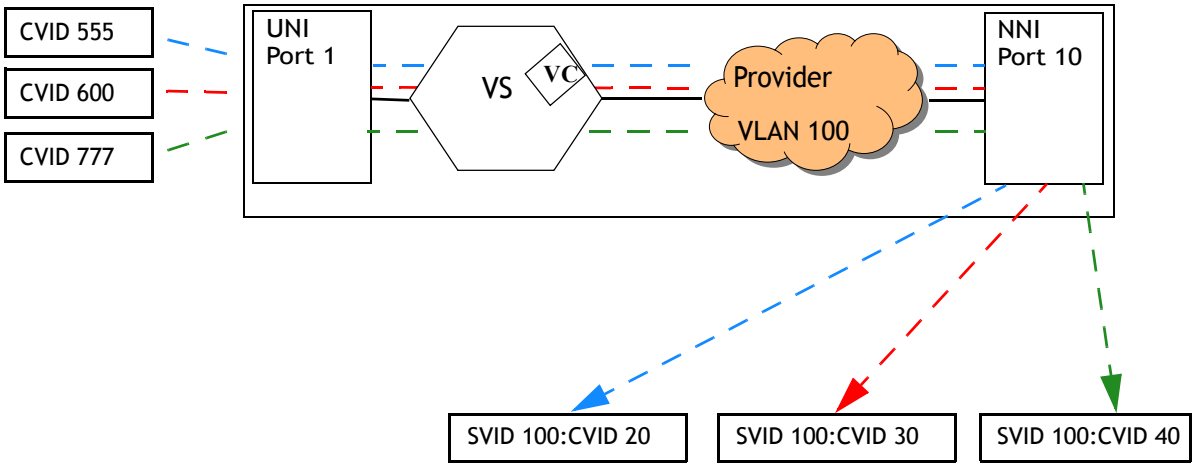
- "When translate-tag is non-zero"
- "When translate-tag is zero"

Copyright© 2012-2013 Ciena® Corporation

When translate-tag is non-zero

In this scenario, frames with CVID 555, 600, or 777 classify to the UNI port 1. Upon ingress, each frame is stamped with the new CVID 20, 30, or 40, respectively, then the SVID 100 is pushed onto it. On egress, frames with the SVID 100 plus the CVID 20, 30, or 40 have the SVID 100 popped off, and the original CVID (555, 600, or 777) is stamped. This configuration requires the allocation of resources as described in "Setting the L2 transform action on a port".

Figure 9-12
VS L2 transform: i-push,e-pop for Q-in-Q



The following example creates the configuration.

```
port set port 1 vs-l2-transform i-stamp:push,e-match-pop:stamp
vlan create vlan 100
vlan add vlan 100 port 10
virtual-circuit ethernet create vc vc100 vlan 100
virtual-switch add reserved-vlan 4000
virtual-switch ethernet create vs vs100 vc vc100 reserved-vlan 4000
virtual-switch ethernet add vs vs100 port UNI-1 vlan 555 translate-tag 20
virtual-switch ethernet add vs vs100 port UNI-1 vlan 600 translate-tag 30
virtual-switch ethernet add vs vs100 port UNI-1 vlan 700 translate-tag 40
```

When translate-tag is zero

In this scenario, the SVID is 100. The Ethernet VC (using provider VLAN 100) is associated with the virtual switch. When that association occurs, the active VLAN changes from the initial "reserved VLAN" to the "provider VLAN". The ports of the virtual switch belong to VLAN 100.

On frame ingress, the device looks up CVID 555 on port UNI-1. On match, CVID 555 is stamped with translate-tag 20 and then pushed with sTag 100.

Copyright© 2012-2013 Ciena® Corporation

On frame egress, the device looks up SVID+(translate tag), that is, <100 + 20>. When the double-tag match is found, the device pops the sTag and stamps the remaining tag with the initial CVID 555.

This process is repeated for each user CVID.

This configuration requires the allocation of resources as described in "Setting the L2 transform action on a port".

The following example creates the configuration.

```
port set port UNI-1 vs-l2-transform i-stamp:push,e-match-pop:stamp
vlan create vlan 100
vlan add vlan 100 port NNI-1
virtual-circuit ethernet create vc X vlan 100
virtual-switch add reserved-vlan 4000
virtual-switch ethernet create vs X vc X reserved-vlan 4000
virtual-switch ethernet add vs X port UNI-1 vlan 555 translate-tag 20
virtual-switch ethernet add vs X port UNI-1 vlan 600 translate-tag 30
virtual-switch ethernet add vs X port UNI-1 vlan 700 translate-tag 40
virtual-switch ethernet add vs X port UNI-1 vlan 800 translate-tag 0
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 9-11
Configuring PFGs

Configure PFGs to restrict forwarding of L2 traffic between groups of ports within a shared VLAN or virtual switch. Note that the PFG feature is not supported on the 3902 switch.

The PFG feature contains three levels of configuration:

- global administrative status, which by default is enabled
- PFG policy set specifications
- VLAN- and virtual switch-based configurations settings

The forwarding restrictions defined by the PFG feature become active, that is, operationally-enabled, when all three levels of configuration are enabled.

By default, the PFG feature is:

- enabled globally
- disabled for all policy sets excluding the management policy
- disabled on all VSs and VLANs, with the exception of the management VLAN which is enabled by default

Copyright© 2012-2013 Ciena® Corporation

Table 9-8 lists supported domains by platform.

Table 9-8
Supported domains by platform

Platform	Domain
3911	management
3916	management, universal
3920	management
3930	management, universal
3931	management, universal
3940	management
3960	management, universal
5140	management
5150	management, universal

Step	Action

1 Disable PFG globally:

```
private-forwarding-groups disable
```

2 Configure PFG port membership:

```
private-forwarding-groups set port <PORT> policy <'mgmt',
'univ'> fwd-group <A,B>
```

where

port <PORT>	is the port.
policy <'mgmt', 'univ'>	is the policy set.
fwd-group <A,B>	is the PFG.

3 Configure the universal policy set:

```
private-forwarding-groups set policy <'mgmt', 'univ'>
[[policy-group-A <A,B>], [policy-group-B <A,B>]]
```

where

policy <'mgmt', 'univ'>	is the policy set.
policy-group-A <A,B>	specifies which interfaces PFG A can forward to.
policy-group-B <A,B>	specifies which interfaces PFG B can forward to.

4 Configure required VLANs and virtual switches:

```
vlan create vlan <VLAN ID>
vlan add vlan <VLAN ID> port <PORT1>
virtual-circuit ethernet create vc <VC ID> vlan <VLAN ID>
virtual-switch add reserved-vlan <VLAN ID2>
virtual-switch ethernet create vs <VS ID> vc <VC ID>
virtual-switch ethernet add vs <VS ID> port <PORT2>
```

5 Enable PFG:

```
private-forwarding-groups enable vlan <VLAN>
private-forwarding-groups enable vs <VS>
private-forwarding-groups enable policy <'mgmt', 'univ'>
private-forwarding-groups enable
```

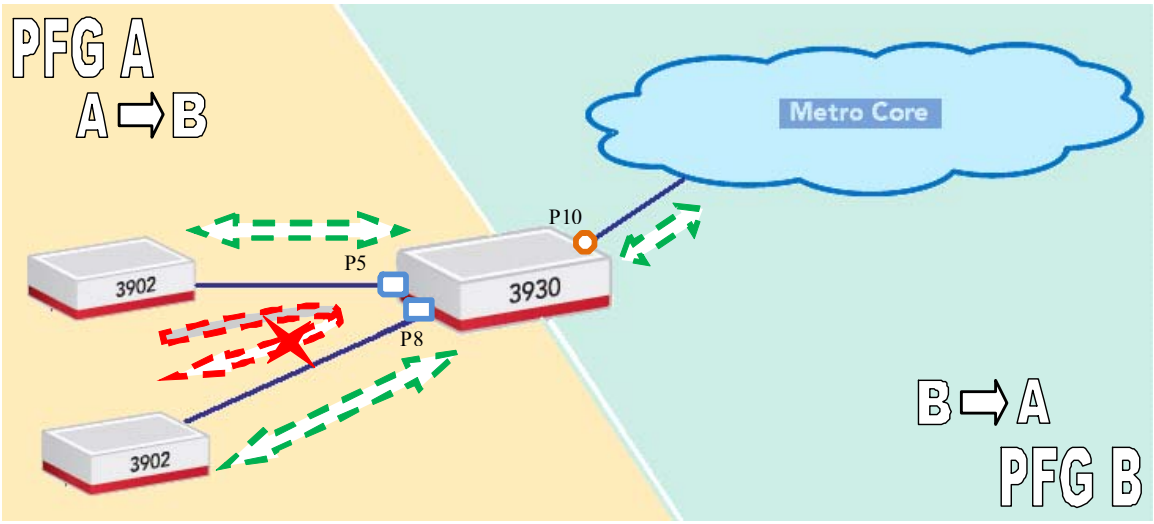
—end—

Example

Figure 9-13 illustrates a simplified example of PFG use on a 3930 platform. Port 10 connects to a core network and ports 5 and 8 are connected to two 3902s. The 3902s represent two separate customer networks.

Copyright© 2012-2013 Ciena® Corporation

Figure 9-13
Sample configuration



The following example configures one service instance while provisioning two attachment customer networks. This configuration sets up a PFG such that the 3902s can access the core network but cannot communicate with each other.

```
private-forwarding-groups disable
private-forwarding-groups set port 10 policy univ fwd-
group B
private-forwarding-groups set port 5 policy univ fwd-
group A
private-forwarding-groups set port 8 policy univ fwd-
group A
private-forwarding-groups set policy univ policy-group-A
B policy-group-B A
vlan create vlan 100
vlan add vlan 100 port 10
virtual-circuit ethernet create vc VC100 vlan 100
virtual-switch add reserved-vlan 1000
virtual-switch ethernet create vs VS1 vc VC100
virtual-switch ethernet add vs VS1 port 5
virtual-switch ethernet add vs VS1 port 8
private-forwarding-groups enable vs VS1
private-forwarding-groups enable policy univ
private-forwarding-groups enable
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 9-12
Disabling the PFG feature

Disable the PFG feature to alter the forwarding restrictions associated with a particular PFG policy set. The PFG feature must be operationally-disabled on all VLANs and virtual switches that reference the policy.

Step	Action
1	Disable the PFG feature on the VLANs that reference the PFG policy set: `private-forwarding-groups disable vlan <VLAN ID>`
2	Disable the PFG feature on the virtual switches that reference the PFG policy set: `private-forwarding-groups disable vs <VS ID>`
3	Disable the PFG policy set: `private-forwarding-groups disable policy <'mgmt', 'univ'>`
4	Disable PFG globally: `private-forwarding-groups disable`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 9-13
Displaying the configuration of EVPL VS members

Display the configuration of EVPL VS members when troubleshooting.

Step	Action
1	Display the configuration of each EVPL virtual-switch member: `virtual-switch ethernet show vs <VirtualSwitch>` where vs <VirtualSwitch> is the virtual switch.

—end—

Example

The following example shows sample output for the virtual switch with an identifier of vs150.

```
> virtual-switch ethernet show vs vs150
+------------- ETHERNET VIRTUAL SWITCH INFO ---------------+
| Parameter             | Value                            |
+-----------------------+----------------------------------+
| Name                  | vs150                            |
| ID                    | 5                                |
| Active VLAN           | (150)   VLAN#150                 |
| Reserved VLAN         | (2503)  VLAN#150                 |
| CPT Status            | Disabled                         |
| CPT Method            | l2pt                             |
| Transparent Validate  | off                              |
| Mac Learning          | Enabled                          |
| VC                    | (1)     vc150                    |
| Subscr Dot1D Policy   | leave                            |
| Decap CoS Policy      | leave                            |
| Decap Fixed Dot1DPri  | 2                                |
| Encap CoS Policy      | fixed                            |
| Encap Fixed Dot1DPri  | 2                                |
+-------------------------------------------------------+
|                 Virtual Switch Members                |
| Port      | Subscriber VLAN | Encap Cos     |Encap Fixed |
|           |                 |  Policy       | Dot1DPri   |
|-----------+-----------------+---------------+------------|
| 5         | 151             | fixed         | 6          |
| 5         | 152             | vs-inherit    | 2          |
+-------------------------------------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 9-14
Displaying virtual switches

Display virtual switches to learn which reserved VLAN is associated with a virtual switch.

Note: The default name for a reserved VLAN cannot be changed. Reserved VLANs always appear in the format `VLAN#XXX`.

Step	Action
1	Display the reserved VLAN associated with a virtual switch: `virtual-switch show`

—end—

Example

The following example shows sample output for the virtual-switch show command.

```
> virtual-switch show

+------------ VIRTUAL SWITCH RESERVED VLANS -----------+
| VLAN                  | Virtual Switch [Attached]     |
+-----------------------+-------------------------------+
| 4094                  | (1)   vs1                     |
+-----------------------+-------------------------------+

+----------------- ETHERNET VIRTUAL SWITCH TABLE -------------------------+
|            |  VLAN  |     |                  | Subscriber |               |
| Name       | Active | CPT | Virtual Circuit  |    Count   |  Description  |
+------------+--------+-----+------------------+------------+---------------+
| vs1        | 4094   | Dis |                  | 0          |               |
+------------+--------+-----+------------------+------------+---------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 9-15
Displaying PFG information

You can display

- the administrative state of the PFG feature
- VLAN-specific PFG settings
- virtual switch-specific PFG settings
- configuration settings for the PFG policy set
- PFG settings for port membership

Step	Action

To display the administrative state of the PFG feature

1 Display the administrative state of the PFG feature:

```
private-forwarding-groups show
```

To display VLAN-specific PFG settings

2 Display VLAN-specific PFG settings:

```
private-forwarding-groups show vlan
```

To display virtual switch-specific PFG settings

3 Display virtual switch-specific PFG settings:

```
private-forwarding-groups show vs
```

To display the configuration settings for the PFG policy set

4 Display the configuration settings for the PFG policy set:

```
private-forwarding-groups show policy
```

To display the PFG settings for port membership

5 Display PFG settings for port membership:

```
private-forwarding-groups show port
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Provider Backbone Bridge Traffic Engineering (PBB-TE) implementation

This chapter details the implementation of Provider Backbone Bridge Traffic Engineering (PBB-TE), formerly referred to as Provider Backbone Transport (PBT), in the core of a Ciena network.

- "VLAN Tagging"
- "Connectivity Fault Management"
- "Benefits"

Note: PBB-TE CLI commands continue to use PBT rather than PBB-TE as the initial implementation was based on PBT.

A basic understanding of the IEEE 802.1Qay, IEEE 802.1Q, 802.1ad, and 802.1ah standards is assumed. In addition, it is highly recommended that CFM (IEEE 802.1ag) be used to monitor the PBB-TE service to warn if any faults occur in the service.

Table 10-1 summarizes PBB-TE for 39XX/51XX switches. Native PBB-TE support is hardware-based.

Table 10-1
PBB-TE support for 39XX/51XX switches

Platform	Non-native PBB-TE	Native PBB-TE	Ports
3902	Not supported	Not supported	Not supported
3911	Not supported	Not supported	Not supported
3916	Not supported	Supported	All ports
3920	Not supported	Not supported	Not supported
3930	Not supported	Supported	All ports
3931	Not supported	Supported	All ports
3940	Supported	Not supported	Enhanced ports only

Copyright© 2012-2013 Ciena® Corporation

Table 10-1
PBB-TE support for 39XX/51XX switches

Platform	Non-native PBB-TE	Native PBB-TE	Ports
3960	Supported	Not supported	Enhanced ports only
5140	Supported	Not supported	Enhanced ports only
5150	Supported (default)	Supported	All ports in native mode Enhanced ports in non-native mode

A software license key is required to activate PBB-TE. Contact Ciena Sales to obtain a license key.

VLAN Tagging

The IEEE 802.1Q standard specifies a mechanism for adding tags to Ethernet frames. This tagging allows an Ethernet network to be divided into virtual networks or VLANs. Generally, individual customers in a provider's network are identified by unique VLAN IDs. Additionally, each customer typically uses VLAN IDs in their own network to differentiate between service types (for example, data or VoIP) and possibly to distinguish between departments. This three-tier hierarchy allows separate domains for the service provider, customer and individual enterprise departments.

Given the 12-bit size of the VLAN field in an Ethernet frame, the number of available VLAN IDs is limited to 4,096. In a large provider network, the number of customers could easily exceed 4,000, exhausting the available VLAN IDs very quickly. To help improve network scalability, two standards have been introduced to support this hierarchical approach, IEEE 802.1ad and IEEE 802.1ah.

Q-in-Q

The IEEE 802.1ad (also known as Q-in-Q, stacked VLANs, or Provider Bridges), extends the original concept of VLAN tags by specifying a new provider-administered 802.1Q tag (Q-tag) that wraps or encapsulates the original customer packet in an additional header. This allows the service provider to administer their own tags to identify individual customer networks, while the first (original) Q-tag is unique within the customer's own network.

This then allows for overlap between the customer and provider VLAN IDs since the customer Q-tag is hidden while it is transported through the provider network. However, while this Q-in-Q approach supports a three-tiered hierarchy and frees up additional VLAN IDs, the service provider can still only

Copyright© 2012-2013 Ciena® Corporation

create 4,094 customer VLANs. Scalability is still an issue as 4,000 tags is simply insufficient for large metropolitan networks. This shortcoming is addressed by the second standard, IEEE 802.1ah.

MAC Header Encapsulation

The IEEE 802.1ah standard, also known as MAC header encapsulation (MAC-in-MAC) or Provider Backbone Bridging, encapsulates the customer MAC header with a MAC header for the service provider. A 24-bit service tag that identifies individual subscribers, is also used in the service provider MAC header. This enables a maximum of 16 million service instances to be supported, thus alleviating the scalability issue entirely.

With IEEE 802.1ah, the overall network is treated as separate service provider and end customer domains. In the service provider domain, the network switches on the service provider MAC header and the customer MAC is not even visible. This introduces strict demarcation between the customer and service provider, enabling a truly hierarchical approach to the network. MAC-in-MAC greatly improves the scalability of Layer 2 networks.

This approach also alleviates another limitation to traditional Ethernet; learning. Traditional Ethernet switch-based networks have two defining characteristics that limit their topological size, namely learning and loop avoidance. When a switch receives a packet destined for an end station it does not yet know about, it floods the packet down every link on which it is connected. This is referred to as Flood for Unknown. As the switch inspects every packet it receives, it remembers or learns the association of sending stations and ingress links. Each switch in the layer 2 domain learns the address and associated link for every other device in the network.

While many metro and core devices may be able to support hundreds of thousands of addresses, requiring each device in a provider's network to handle this number of addresses is cost prohibitive and impacts protection switching schemes. MAC-in-MAC encapsulation solves this Layer 2 scalability problem since there are now fewer MAC addresses to learn at the core.

The IEEE 802.1ah MAC header, specifies that a backbone VLAN tag (B-tag), and an instance service VLAN tag (I-tag) be added to packets traversing the provider backbone bridge. This extended service VLAN tag is mapped from the Service VLAN ID tag and is not any longer than a normal VLAN tag. The additional B-tag is then added to ensure that the switches in the middle of the 802.1ah core do not need to know about the 802.1ah functions, thus ensuring backward compatibility. The backbone MAC header is then removed at the other end of the provider's 802.1ah backbone bridge. A big advantage of 802.1ah is that it works in conjunction with the 802.1ad Provider Bridges VLAN stacking standard to allow both techniques to be used simultaneously.

Copyright© 2012-2013 Ciena® Corporation

PBB-TE Tunnels

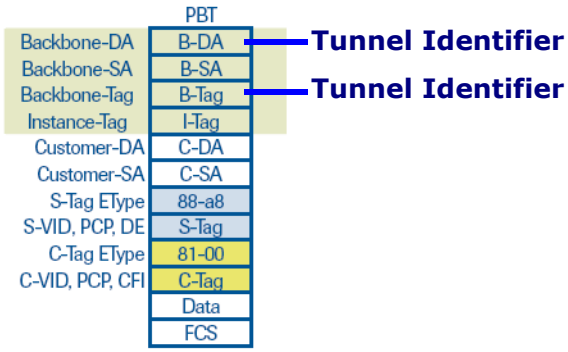
Note: When using PBB-TE, RSTP/MSTP should not be enabled as all PBB-TE ports need to be in the forwarding state. For example, if RSTP/MSTP is enabled it could block one of the PBB-TE ports and thus remove the backup tunnel.

PBB-TE has emerged to further address limitations related to scalability and reliability. PBB-TE may be deployed in place of PBB or may run in parallel with PBB. In both cases, PBB-TE eliminates the need for backbone core devices to perform learning and flooding. Instead, point-to-point tunnels to transport L2VPNs are provisioned. Rather than using RSTP/MSTP to prevent loops, the tunnels are used to traffic engineer the provider backbone network.

A PBB-TE tunnel is unidirectional and is set up to ingress on a local node at one edge of the provider backbone and to terminate on a remote device at the other end of the backbone. This tunnel transports encapsulated subscriber frames from the local node to the remote node. A second tunnel is then configured in the opposite direction to ingress the remote node and terminate on the local node. This serves as the transport for the reverse traffic. A tunnel pair together make a bidirectional pipe referred to as a PBB-TE trunk.

PBB-TE tunnels are set up to transport Ethernet Virtual Circuits (IEEE 802.1ad), with the subscriber frames encapsulated using 802.1ah MAC-in-MAC encapsulation. Tunnels are then identified by the Backbone Destination Address (B-DA) and the Backbone Tag (B-Tag). In Figure 10-1, the highlighted fields are added to the frame to identify tunnels and to identify flood domains and interconnects.

Figure 10-1
PBB-TE Frame Format

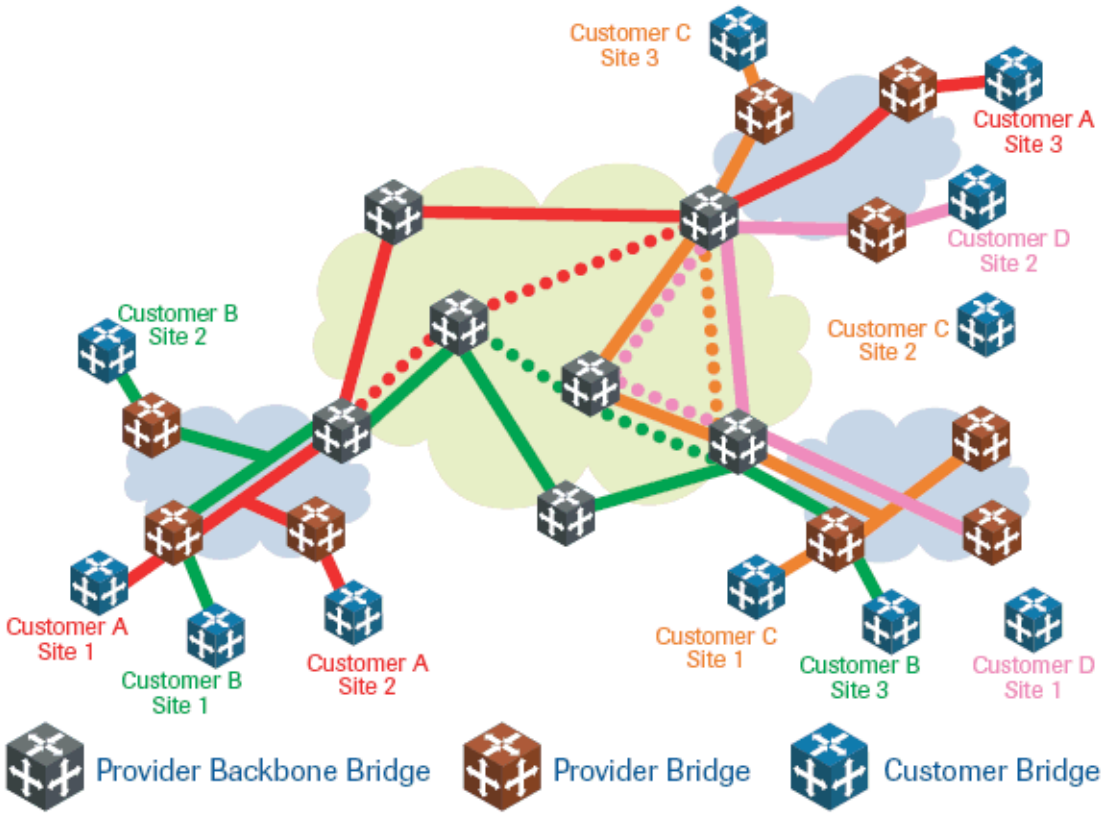


Since the tunnels are point-to-point, PBB-TE can also achieve recovery times approaching 50 ms. Providers can set up primary and protection PBB-TE paths and then leverage Carrier Ethernet OAM standard CFM (IEEE 802.1ag) to monitor the tunnels. This provides fault notifications in milliseconds and thus carrier-grade failover times can be achieved.

Copyright© 2012-2013 Ciena® Corporation

Figure 10-2 shows an example PBB-TE network with the core provider network in the center and three access networks attached to the core. In the core, primary tunnels are configured (solid lines) as well as backup tunnels (dotted lines) in order to provide resiliency and greater utilization of the backbone network. Again, in the core, there is no loop detection and no learning.

Figure 10-2
Example of a PBB-TE Network



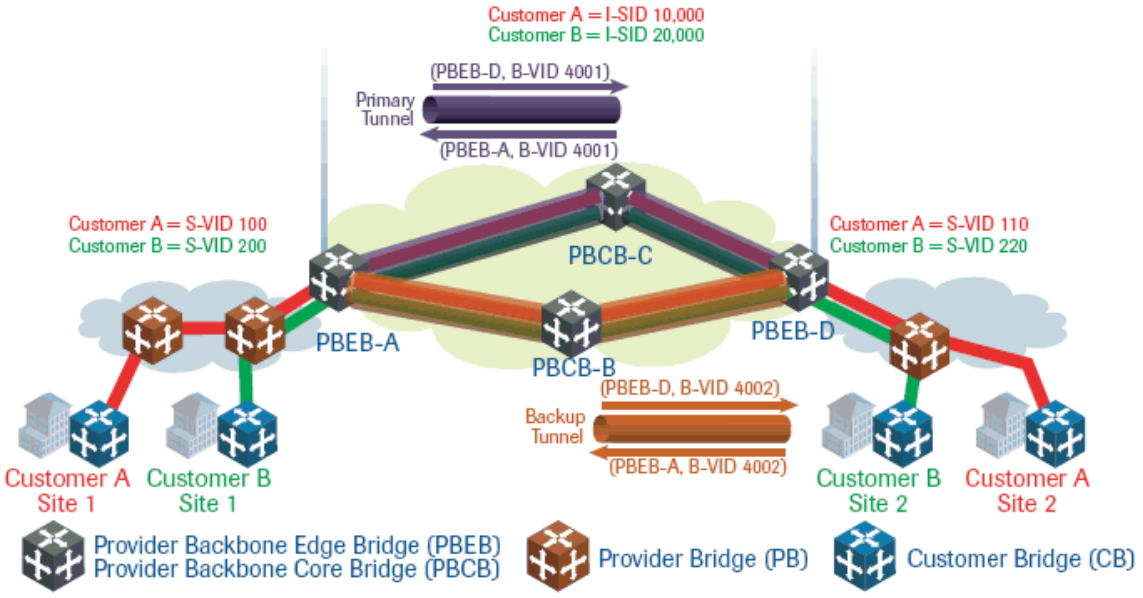
The system software supports configuration of a maximum of 32 PBB-TE bidirectional tunnels. However, one PBB-TE tunnel may be used to transport multiple subscriber services. In addition, one virtual circuit must be created for each service instance. A maximum of 127 virtual circuits are supported by one PBB-TE tunnel. PBB-TE tunnels may also be configured across link aggregation ports.

The PBB-TE virtual circuit must be attached to a virtual switch before subscriber ports are attached. When detaching a PBB-TE virtual circuit from a virtual switch, subscriber ports must be detached first.

Note: In PBB-TE native mode, PBB-TE subscribers and Q-in-Q subscribers cannot be on the same port.

Figure 10-3 depicts two Provider Bridged networks interconnected by a PBB-TE network. Two customer Layer 2 Virtual Private networks are shown traversing primary and backup PBB-TE tunnels through the core network.

Figure 10-3
A PBB-TE Network with Primary and Backup Tunnels



Traffic from Customer A originates at its Site 1. The Provider Bridge encapsulates the customer traffic by adding a Subscriber S-Tag containing the configured S-VID value of 100 reserved for Customer A within its domain. This traffic is then sent to the Provider Backbone Edge Bridge A (PBEB-A).

PBEB-A has been configured to assign Customer A traffic (S-VID=100) to a 24-bit Instance Service Identifier (I-SID) value of 10,000. This I-SID value is associated with a primary and a backup PBB-TE tunnel. Each primary and backup tunnel is identified using the combination of a PBEB Destination MAC address and a Backbone-VID (B-VID). This is a significant difference between PBB-TE and PBB. In PBB, B-VIDs represent flood domains that interconnect multiple Provider Bridged networks. In PBB-TE, B-VIDs together with B-DAs identify which tunnel the frame will use.

In our example, the PBEB-A encapsulates S-VID 100 traffic by adding the B-DA value for PBEB-D, a B-SA value for PBEB-A, a B-VID value of 4001 (primary tunnel), and the I-SID value of 10,000. This MAC header encapsulated traffic is forwarded to Provider Backbone Core Bridge-C (PBCB-C). PBCB-C has been configured to not learn or flood traffic on B-VID

Copyright© 2012-2013 Ciena® Corporation

4001, which has been reserved for PBB-TE use. The fact that PBB-TE does not learn or flood is an important point. Each PBCB device must be provisioned with forwarding database entries in order to properly forward traffic within the tunnels.

The PBCB-C forwarding table learns an entry for the combination of PBEB-D, B-VID 4001 and the traffic is forwarded on the particular port in the direction of PBEB-D. PBEB-D receives the traffic and removes the MAC header encapsulation. Since the S-VID values are only locally significant to the Provider Bridged network, a provider has the flexibility to translate the S-VID value. In our example, PBEB-D has been configured to associate I-SID 10,000 with S-VID 110.

In this illustration, traffic from the tunnel is de-encapsulated and the S-VID is remapped to the value of 110. The traffic is forwarded to the provider bridge attached to Customer A's Site 2. The S-Tag encapsulation is removed by the PB device and the original customer frame from Site 1 is delivered to Site 2.

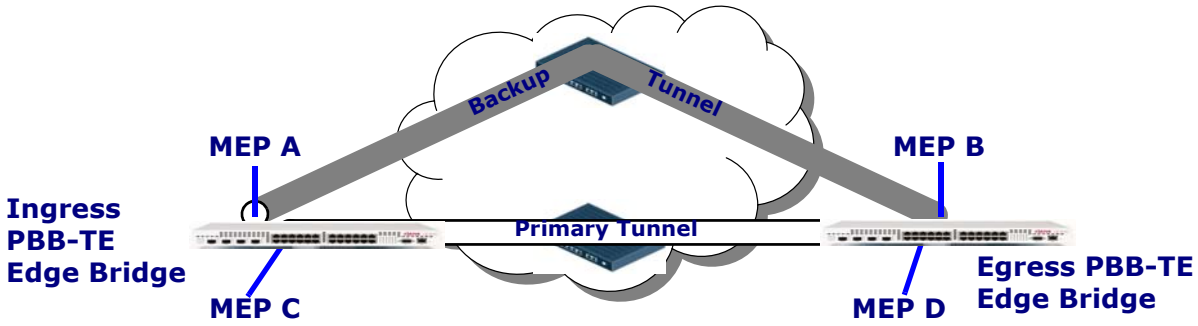
> ***Note:*** It should also be noted that the PBB-TE service tag is 22 Bytes in length and this needs to be taken into account when calculating available bandwidth. For example, when sending a 64-Byte packet at 100% of line rate through a PBB-TE tunnel, each 64-Byte frame is actually 86 Bytes long (64 + 22).

Connectivity Fault Management

PBB-TE virtual switches and tunnels are monitored through the use of IEEE 802.1ag Connectivity Fault Management (CFM) Continuity Check Messages (CCM). When CFM services are configured for PBB-TE virtual switches, CCM control frames are sent and received within the PBB-TE tunnels and monitor faults based on the associated PBB-TE virtual circuit events. For CFM services that monitor PBB-TE tunnels, CCM control frames are sent and received regularly along the PBB-TE tunnels. If the primary tunnel should experience a fault, the tunnel endpoints automatically begin using the backup tunnel. The forwarding database entries are pre-configured along the backup path to minimize the failover and restoration times. Only one backup tunnel can be created.

Copyright© 2012-2013 Ciena® Corporation

Figure 10-4
PBB-TE with CFM



Note: Inner tag priority is used for scheduling frames when service levels and mappings are applied to ports associated with PBB-TE tunnels. To ensure proper QoS treatment of transiting CCM frames, which have a single B-VID tag, an outer tag is pushed onto the frame. The addition of this outer tag enables the inner tag priority to be detected correctly so that these CCM frames are scheduled properly. Upon egress, the outer tag is popped.

It should be noted that once a failover occurs (switching the path from the primary to the backup tunnel), the default behavior is for the provisioned primary tunnel to become the backup tunnel unless tunnel reversion is configured using the `pbt set tunnel-reversion` command. The tunnel-reversion command will automatically revert traffic back to the provisioned primary tunnel once it is up and the reversion hold timer is completed (`pbt set reversion-hold-time`).

Example

```
> pbt set tunnel-reversion off
> pbt set reversion-hold-time 3000
```

Manual reversion allows the network operator to control when or if the reversion will take place (such as late at night or on a weekend) so that the impact to the service is minimum. Manual reversion also prevents flapping between the two tunnels if the primary is not entirely stable. Automatic tunnel reversion is useful if the backup tunnel is inferior to the primary tunnel. Thus, failback occurs automatically.

Note: When configuring PBB-TE reversion, tunnel reversion must be enabled on all remote bridges and the reversion hold time must be set to the same value. If reversion settings are not symmetrical, it can cause issues with indicating the primary/backup tunnel and can cause delays in reversions.

Copyright© 2012-2013 Ciena® Corporation

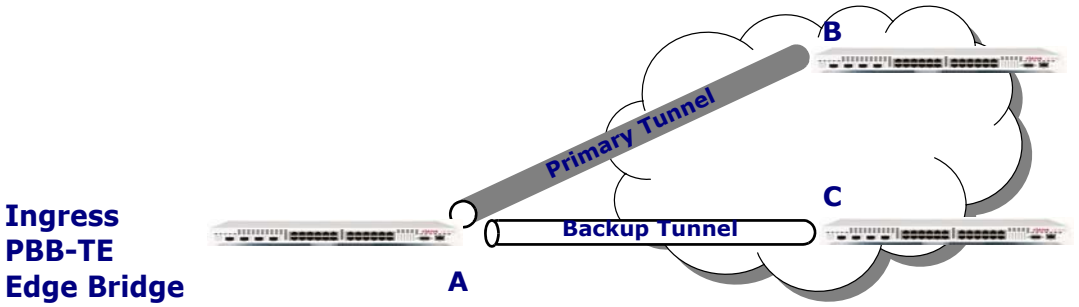
Ciena supports the draft IEEE 802.1ag CFM including MAC ping, MAC traceroute, and CCM. Using these powerful CFM tools, both service connectivity and PBB-TE tunnel resiliency is enhanced. In addition, CCM thresholds can be configured to fine-tune protection switching.

PBB-TE Dual Homing

The system software also supports the dual homing of primary and backup tunnels. This enables PBB-TE tunnels to terminate on entirely separate devices. This offers device redundancy and path diversity for upstream connections.

In Figure 10-5, device A is dual homed to device B and device C. This offers link resiliency, such that if device B fails, traffic is then diverted to the backup tunnel terminating at device C.

Figure 10-5
PBB-TE Dual Homing



Dual homing is configured by specifying a unique MAC address and bridge name when using the `pbt remote-bridge create bridge-name` command. This second remote bridge is then specified when creating the backup tunnel.

Tunnel Pairing and Synchronization

Although PBB-TE tunnels are set up as trunks, and both an encapsulation and a decapsulation tunnel are required for traffic transport, the tunnel states are not automatically linked together. This makes it possible for the encap and decap tunnels to have different operational states.

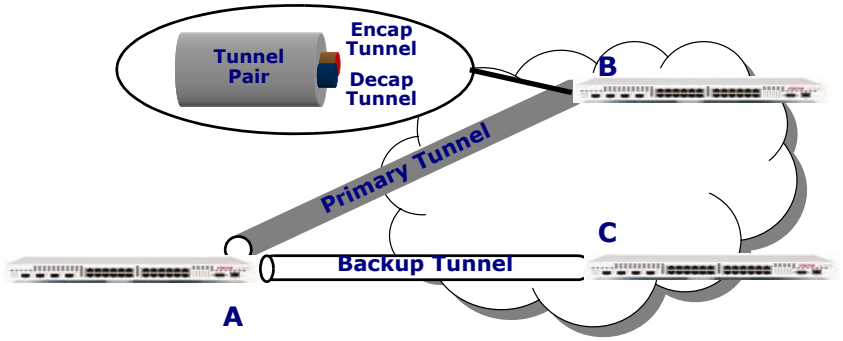
For example, if the encap tunnel was down due to misconfiguration, the decap tunnel could still be up and running. This would allow incoming traffic to be decapsulated, thus providing only unidirectional traffic on the circuit. This creates a problem with devices on the terminating end of the tunnels. They would accept all frames regardless of their B-VID value.

Copyright© 2012-2013 Ciena® Corporation

To avoid this behavior, encap and decap tunnels should be tied together using the `tunnel pair create` command. By pairing encap and decap tunnels, traffic on an interface is dropped unless the B-VID is specifically configured. In addition, it synchronizes the operational state of the two tunnels. When one tunnel goes down, encap or decap, CCMs are sent over the interface to ensure that the device on the other end takes both tunnels down. The tunnel state will remain down until the MEPs on the other devices have verified connectivity.

In Figure 10-6, the encap tunnel on device B goes down. Device B then sends a CCM error code to device A. Device A then takes down both the primary encap and decap tunnels and a failover occurs to the backup tunnel.

Figure 10-6
Tunnel Pairing Example



To ensure synchronization between primary and backup tunnels, you can enable tunnel synchronization. When tunnel synchronization is turned on for the primary tunnel, a deactivation of either the primary or backup tunnel causes CCMs to be transmitted to the inactive tunnel. This feature enables interoperability with devices running the same system software. By default, tunnel synchronization is disabled.

Benefits

The main benefits of PBB-TE include:

- Removing the 4,000 tag limitation, enabling 16 million distinct services to be configured.
- No learning or flooding in the core of the network for a reduction in complexity and cost.
- Customer MAC address and other information is tunneled through the provider's network, enhancing security and scalability.
- Using specifically engineered paths or tunnels allows the provider to target maximum utilization of the core network devices.

Copyright© 2012-2013 Ciena® Corporation

- The customer and Provider control domains are separated, allowing layer 2 control frames to be transported through the provider's network.
- 802.1ag CFM can be used to monitor tunnels and provide carrier-grade failover detection.

This chapter provides the following procedures for PBB-TE:

- "Verifying that a port can participate in PBB-TE" on page 10-12
- "Switching from non-native to native PBB-TE support" on page 10-13
- "Enabling tunnel synchronization" on page 10-14
- "Disabling tunnel synchronization" on page 10-15
- "Configuring PBB-TE" on page 10-16
- "Releasing reserved BVIDs" on page 10-34
- "Displaying PBB-TE information" on page 10-35

Copyright© 2012-2013 Ciena® Corporation

Procedure 10-1
Verifying that a port can participate in PBB-TE

Verify that a port can participate in PBB-TE before using the port in a PBB-TE configuration. Ports that can participate in PBB-TE have enhanced capabilities.

You can display the capabilities of

- a selected port
- all ports
- ports on a blade

Note: All ports on the 3916, 3930, 3931, and 5150 platforms are supported by native PBB-TE.

Step	Action

To display the capabilities of a selected port

1 Display the capabilities of a selected port:

```
port show port <port> capabilities
```

where

port <port>	is the port that you want to verify for PBB-TE participation.
capabilities	displays the capabilities of the port.

To display the capabilities of all ports

2 Display the capabilities of all ports:

```
port show capabilities
```

To display the capabilities of all ports on a blade

3 Display the capabilities of all ports on a blade:

```
blade show capabilities
```

—end—

Procedure 10-2
Switching from non-native to native PBB-TE support

The 5150 platform can operate in non-native mode or native mode for PBB-TE support. Native mode is hardware-based. By default, the 5150 platform operates in non-native mode.

In non-native mode, the 5150 supports PBB-TE on enhanced ports only.

In native mode, the 5150 platform supports PBB-TE on all ports.



CAUTION
No automated conversion of management configuration between non-native mode and native mode on 5150 platform
If you are currently using management over PBB-TE and you plan to switch from non-native mode to native mode on a 5150 platform, contact Ciena support personnel to obtain details on how this conversion is best accomplished.

Step	Action
1	Switch from non-native to native PBB-TE support: `pbt set mode native`
2	Save the configuration: `configuration save`
3	Reboot the device: `chassis reboot now`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 10-3
Enabling tunnel synchronization

You can enable tunnel synchronization

- when creating a tunnel
- for an existing tunnel

Step	Action

To enable tunnel synchronization when creating a tunnel

1 Enable tunnel synchronization when creating a tunnel:

```
tunnel encap create static-pbt <static-pbt> tunnel-sync
on
```

where

<static-pbt> is the name of the tunnel.

To enable tunnel synchronization for an existing tunnel:

2 Enable tunnel synchronization for an existing tunnel:

```
tunnel encap set static-pbt <static-pbt> tunnel-sync on
```

where

<static-pbt> is the name of the tunnel.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 10-4
Disabling tunnel synchronization

You can disable tunnel synchronization

- when creating a tunnel
- for an existing tunnel

Step	Action

To disable tunnel synchronization when creating a tunnel

1 Disable tunnel synchronization when creating a tunnel:

```
tunnel encap create static-pbt <static-pbt> tunnel-sync
off
```

where

<static-pbt> is the name of the tunnel.

To disable tunnel synchronization for an existing tunnel:

2 Disable tunnel synchronization for an existing tunnel:

```
tunnel encap set static-pbt <static-pbt> tunnel-sync off
```

where

<static-pbt> is the name of the tunnel.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 10-5
Configuring PBB-TE

Configuring PBB-TE comprises the following:

- Prepare the device for PBB-TE
- Configure PBB-TE parameters
- Configure PBB-TE tunnels
- Optionally configure PBB-TE back-up tunnels
- Configure the virtual circuit
- Configure the virtual switch
- Configure in-band management (non-native PBB-TE)
- Configure in-band management (native PBB-TE platform
- Configure CFM on a PBB-TE virtual switch
- Configure CFM on a PBB-TE tunnel
- Optionally configure QoS on a PBB-TE tunnel

To use PBB-TE, a license key is required in order to enable PBB-TE features. PBB-TE license keys can be purchased by contacting Ciena sales.

> ***Note:*** It is recommended that you manually set the B-SA PBB-TE bridge MAC address. If you do not set the B-SA PBB-TE bridge MAC address, the system software will send the chassis base MAC address as the B-SA. In the event of a hardware failure, the chassis base MAC would change, requiring configuration changes on other devices configured to use this device as a remote bridge. This setting needs to be unique from the B-SA for all devices.

Note that configuration of the backup tunnel is optional, but recommended. The tunnel name must also be the same as the primary tunnel.

Configuring a virtual circuit for PBB-TE involves setting the ingress I-SID for service identification, the egress I-SID, the destination node, and the PBB-TE tunnels to be used.

> ***Note:*** The system can automatically assign a reserved VLAN to a PBB-TE virtual switch upon creation or you can optionally specify a reserved VLAN. This reserved VLAN becomes the S-VID in the S-Tag used upon encapsulation. You can configure the VC to retain this S-Tag upon encapsulation by setting the VC `retain-stag` attribute to `yes`. When this

Copyright© 2012-2013 Ciena® Corporation

attribute is set to yes, then the decapsulated frame must have the S-Tag that matches the reserved VLAN of the VS and the egress ISID of the VC; otherwise, it will be dropped.

The number of VCs supported for PBB-TE depends upon the platform capabilities as shown in Table 10-2.

Table 10-2
PBB-TE virtual circuits supported per platform

Platform	VCs
3916,3930,3931	256
3940, 5140	128
3960, 5150	512

Note: The number of VCs supported for PBB-TE is reduced by the number of VCs configured for Q-in-Q.

Configuring a virtual switch for PBB-TE involves creating reserved VLANs, designating the PBB-TE VC to be used with the virtual switch, and adding subscriber ports/VLANs. Optionally, Encap CoS policies may also be set for CoS mapping between the customer frames and the encapsulation PBB-TE tunnel header.

The number of VSs and members supported for PBB-TE depends upon the platform capabilities as shown in Table 10-3.

Table 10-3
PBB-TE virtual switch capabilities

Platform	VSs	VS members per port	VS members per switch
3916,3930,3931	256	512	512
3940, 5140	128	512	512
3960	512	2048	2048
5150	512	3584	3584

Note: For 3940, 3960, 5140, and 5150 platforms, the number of VSs and VS members per port and switch supported for PBB-TE are reduced by the number of VSs and VS members per port and switch configured for Q-in-Q.

Copyright© 2012-2013 Ciena® Corporation

Management VLAN traffic can be tunneled through a PBB-TE tunnel in non-native mode by creating a management VC. The management VC creation implicitly creates a management VS that can be displayed using `show` commands but there are not any operations that can be performed on the management VS. The management VC (mgmt-vc) will automatically take the place of the remote management interface vlan. Only one management VC can be created, and ports must be explicitly added to and deleted from the management VLAN in order to be encapsulated in a PBB-TE tunnel.

When using native mode, management traffic can be tunneled through a PBB-TE tunnel by using a management VS. Such a management VS is manually configured and explicitly associated with the remote interface by the user.

Configuring CFM for the PBB-TE virtual switch requires enabling CFM globally, creating a CFM service using the PBB-TE virtual switch, and then enabling the service. CCMs and other CFM frames are transmitted out the PBB-TE virtual circuit as PBB-TE data frames with full PBB-TE service encapsulation and are forward to the virtual switch member ports after the PBB-TE service header is removed.

Configuring CFM for the PBB-TE tunnels requires enabling CFM globally, creating a CFM service using the PBB-TE tunnel, and then enabling the service.

Configuring QoS for PBB-TE involves setting any traffic metering required at the ingress (UNI) and setting the CoS marking policy. Only advanced mode is supported on platforms with native PBB-TE, that is, 3916, 3930, 3931, and 5150.

For additional virtual switch CoS policies, see the Command Reference Manual.

Step	Action

To prepare the device for PBB-TE

1 View software licence information:

```
software license show
```

2 Ensure that MSTP or RSTP is disabled either globally or on the individual encap and decap tunnel port:

```
mstp disable
mstp show
```

or

```
rstp disable
rstp show
```


Copyright© 2012-2013 Ciena® Corporation

To configure PBB-TE parameters

3 Configure the PBB-TE bridge MAC that will be used as the B-SA in the frames.

```
pbt set bridge-mac <MAC address: XX:XX:XX:XX:XX:XX>
reversion-hold-time <MILLISECONDS> service-tag-ethertype
<NUMBER: 1-65535> transit-tunnel-ethertype-remark <on |
off> tunnel-reversion <on | off> tunnel-tag-ethertype
<NUMBER: 1-65535> pbt set stag-tpid <8100|9100|88A8> mode
<native|non-native>
```

where

bridge-mac <MAC address: XX:XX:XX:XX:XX:XX>	is the MAC address for the device used for PBT.
reversion-hold-time <MILLISECONDS>	is reversion hold time in milliseconds.
service-tag-ethertype <NUMBER: 1-65535>	is the ether-type for the I-tag. In native mode, the default value of 0x88e7is used: it is not configurable. ***Note:*** to avoid potential traffic loss, do not use Ethertype values of 0x8100, 0x88a8, or the configured CFM Ethertype (default 0x8902). In addition, switches are hard-coded to recognize 0x88e7 as a valid service tag. PBB-TE frames with the 0x88e7 service tag are received and processed in addition to the configured service tag. In order to transmit PBB-TE frames with the I-TAG of 0x88e7, set the service tag to 0x88e7.
transit-tunnel-ethertype-remark <on | off>	is the remark B-tag ether-type for frames in transit.
tunnel-reversion <on | off>	sets the reversion behavior from backup to primary.

Copyright© 2012-2013 Ciena® Corporation

where

tunnel-tag-ethertype <NUMBER: 1-65535>	is the ether-type for the B-tag. The default value is 0x88a8. The etype values are restricted to 0x8100, 0x9100 and 0x88a8. ***Note:*** In order for PBB-TE frames to be transmitted through a legacy VLAN switch, the PBB-TE tunnel tag Ethertype must be set to 0x8100 (33024 decimal). This is not supported in native PBB-TE (3916,3930,3931, and 5150), and therefore transit switches must be configured to mark packets as 88a8 such as: `vlan set vlan 100 egress-tpid 88a8 virtual-circuit ethernet set port 7,10 vlan-ethertype-policy vlan-tpid`
pbt set stag-tpid <8100|9100|88A8>	is the stag-tpid etype. This attribute is not configurable in non-native mode.
mode <native| non-native>	is the PBB-TE mode.

4 Reserve Bridge VLAN IDs (B-VIDs). Each incoming tunnel requires a unique B-VID.

```
pbt reserve bvid <VLAN ID List>
```

where

bvid <VLAN ID List>	is the VLAN list. ***Note:*** The maximum number of B-VIDs that can be reserved is equal to the maximum number of PBB-TE tunnels, which is 32.

5 Configure remote bridge names for use in tunnels and virtual circuits. For operator convenience, tunnel endpoints can be identified with names by associating a MAC address with name. The subsequent Tunnel and VC configurations refer to the nodes by these names.

```
pbt remote-bridge create bridge-name <bridge-name>
bridge-mac <MAC address: XX:XX:XX:XX:XX:XX>
```

where

bridge-name <bridge-name>	is the bridge name.
bridge-mac <MAC address: XX:XX:XX:XX:XX :XX>	is the MAC address for the node.

Copyright© 2012-2013 Ciena® Corporation

To configure PBB-TE tunnels

6 Configure the encapsulation tunnel:

```
tunnel encap create static-pbt <static-pbt> {b-vid
<NUMBER: 2-4094>} {dest-bridge-name <PBT Bridge Name>}
encap-cos-policy <fixed | vc-inherit | vs-inherit> encap-
fixed-pcp <NUMBER: 0-7> {port <Port>} tunnel-sync
<on|off>
```

where

static-pbt <static-pbt>	is the primary tunnel name.
b-vid <NUMBER: 2-4094>	is the encapsulation B-VID.
dest-bridge-name <PBT Bridge Name>	is the destination bridge.
encap-cos-policy <fixed | vc-inherit | vs-inherit>	is the encapsulation CoS policy.
encap-fixed-pcp <NUMBER: 0-7>	is the encapsulation fixed PCP.
port <Port>	is the outgoing port for the tunnel.
tunnel-sync <on|off>	is the state of tunnel synchronization.

7 Configure the decapsulation tunnel:

```
tunnel decap create static-pbt <static-pbt> {b-vid
<NUMBER: 2-4094>} {port <Port>} {src-bridge-name <PBT
Bridge Name>}
```

where

static-pbt <static-pbt>	is the tunnel name.
b-vid <NUMBER: 2-4094>	is the decapsulation B-VID.
port <Port>	is the incoming port for the tunnel.
src-bridge-name <PBT Bridge Name>	is the source bridge.

Copyright© 2012-2013 Ciena® Corporation

8 Pair the primary encapsulation and decapsulation tunnels:

```
tunnel pair create tnl-pair <tnl-pair> {encap-pbt <PBT
Primary Tunnel Name>] [encap-backup-pbt <PBT Backup
Tunnel Name>] {decap-pbt <PBT Decap Tunnel Name>}
```

where

tnl-pair <tnl-pair>	is the tunnel pair name.
encap-pbt <PBT Primary Tunnel Name>	is the primary encapsulation tunnel. ***Note:*** The BVID must be the same for paired encapsulation and decapsulation tunnels.
encap-backup-pbt <PBT Backup Tunnel Name>	is the backup encapsulation tunnel.
decap-pbt <PBT Decap Tunnel Name>	is the decapsulation tunnel.

To configure the PBB-TE backup tunnels

9 Configure the backup encapsulation tunnel:

```
tunnel encap create static-backup-pbt <static-backup-pbt>
{b-vid <NUMBER: 2-4094>} {dest-bridge-name <PBT Bridge
Name>} encap-cos-policy <fixed | vc-inherit | vs-inherit>
encap-fixed-pcp <NUMBER: 0-7> {port <Port>}
```

where

static-backup-pbt <static-backup-pbt>	is the backup tunnel name.
b-vid <NUMBER: 2-4094>	is the encapsulation B-VID.
dest-bridge-name <PBT Bridge Name>	is the destination bridge.
encap-cos-policy <fixed | vc-inherit | vs-inherit>	is the encapsulation CoS policy.
encap-fixed-pcp <NUMBER: 0-7>	is the encapsulation fixed PCP.
port <Port>	is the outgoing port for the tunnel.

Copyright© 2012-2013 Ciena® Corporation

10 Configure the backup decapsulation tunnel:

```
tunnel decap create static-pbt <static-pbt> {b-vid
<NUMBER: 2-4094>} port <Port> src-bridge-name <PBT Bridge
Name>}
```

where

static-pbt <static-pbt>	is the tunnel name.
b-vid <NUMBER: 2-4094>	is the decapsulation B-VID.
port <Port>	is the incoming port for the tunnel.
src-bridge-name <PBT Bridge Name>	is the source bridge.

11 Pair the backup encapsulation and decapsulation tunnels:

```
tunnel pair create tnl-pair <tnl-pair> encap-backup-pbt
<PBT Backup Tunnel Name> decap-pbt <PBT Decap Tunnel Name>
```

where

tnl-pair <tnl-pair>	is the tunnel pair name.
encap-backup-pbt <PBT Backup Tunnel Name>	is the backup encapsulation tunnel. ***Note:*** The BVID must be the same for paired encapsulation and decapsulation tunnels.
decap-pbt <PBT Decap Tunnel Name>	is the decapsulation tunnel.

12 Display the tunnel configuration information to verify the configuration:

```
tunnel encap show
```


Copyright© 2012-2013 Ciena® Corporation

To configure the virtual circuit

13 Configure the virtual circuit:

```
virtual-circuit pbt create static-vc <static-vc> {egress-isid <NUMBER: 1-16777215>} {ingress-isid <NUMBER: 1-16777215>} {mgmt-vc <no|yes>} {tunnel <PBT Primary Tunnel Name>} {retain-stag <no|yes>} {s-tag <NUMBER: 1-4095>}
```

where

static-vc <static-vc>	is the PBB-TE virtual circuit name.
egress-isid <NUMBER: 1-16777215>	is the egress I-SID. ***Note:*** Egress I-SID and ingress I-SID values must be the same when using native PBB-TE, that is, 3916, 3930, 3931, and 5150 platforms.
ingress-isid <NUMBER: 1-16777215>	is the ingress I-SID.
mgmt-vc <no | yes>	sets the virtual circuit to be a management virtual circuit.
tunnel <PBT Primary Tunnel Name>	is the encapsulation PBB-TE tunnel name.
retain-stag <no | yes>	indicates whether to retain Q-in-Q tag over PBB-TE encapsulation/decapsulation. If the virtual switch is to be used for in-band management in native PBB-TE mode, set this attribute to no.
s-tag <NUMBER: 1-4095>	is the Q-in-Q tag value. If the retain-stag attribute is set to yes, native PBB-TE mode requires that an s-tag value be specified explicitly.

To configure the virtual switch

14 Reserve VLANs for the virtual switch:

```
virtual-switch add reserved-vlan <VLAN ID List>
```

where

reserved-vlan <VLAN ID List>	is the VLAN list.

Note: The system can automatically assign the reserved VLAN from this pool when the virtual switch is created or you can optionally specify a reserved VLAN when you create the virtual switch by setting the `reserved-vlan` attribute.

Copyright© 2012-2013 Ciena® Corporation

15 Create the virtual switch naming it and associating it with the PBB-TE virtual circuit:

```
virtual-switch ethernet create vs <vs> reserved-vlan
<NUMBER: 1-4094> vc <Virtual Circuit Name>
```

where

vs <vs>	is the virtual switch name.
reserved-vlan <NUMBER: 1-4094>	is the reserved VLAN.
vc <Virtual Circuit Name>	is the Ethernet virtual circuit.

16 Add the subscriber port and VLAN to the virtual switch:

```
virtual-switch ethernet add vs <vs> {port <port>} {vlan
<VLAN list>} {ip-interface <interface-name>} {encap-cos-
policy <dot1dpri-inherit | fixed | ip-prec-inherit |
phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-
dot1dpri <NUMBER: 0-7>} [statistics <on | off>]
[translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100
| 9100 | 88A8>}
```


Copyright© 2012-2013 Ciena® Corporation

To configure in-band management (non-native PBB-TE)

17 Create a management virtual circuit:

```
virtual-circuit pbt create static-vc <static-vc> {egress-isid <NUMBER: 1-16777215>} {ingress-isid <NUMBER: 1-16777215>} mgmt-vc <no | yes> tunnel <PBT Primary Tunnel Name> retain-stag <no | yes> {s-tag <NUMBER: 1-4095>}
```

where

static-vc <static-vc>	is the PBB-TE virtual circuit name.
egress-isid <NUMBER: 1-16777215>	is the egress I-SID. ***Note:*** Egress I-SID and ingress I-SID values must be the same when using native PBB-TE, that is, 3916, 3930, 3931, and 5150 platforms.
ingress-isid <NUMBER: 1-16777215>	is the ingress I-SID.
mgmt-vc <no | yes>	sets the virtual circuit to be a management virtual circuit.
tunnel <PBT Primary Tunnel Name>	is the encapsulation PBB-TE tunnel name.
retain-stag <no | yes>	indicates whether to retain Q-in-Q tag over PBB-TE encapsulation/decapsulation.
s-tag <NUMBER: 1-4095>	is the Q-in-Q tag value.

Note: Note: IPv6 in-band management is supported over native PBB-TE.

To configure in-band management (native PBB-TE platform)

18 Set the tag-vid for the VS created in step 15 to enable it to handle management traffic:

```
virtual-switch cpu-if set vs <vs> tag-vid <NUMBER: 1-4094>
```

where

vs <vs>	is the virtual switch name.
tag-vid <NUMBER: 1-4094>	is the inner tag for PBB-TE management frames. The default customer management VLAN is 127.

Note: IPv6 in-band management is not supported over non-native PBB-TE.

Copyright© 2012-2013 Ciena® Corporation

19 Set interface attributes:

```
interface remote set vs <vs>
```

where

vs <vs> is the management virtual switch.

To configure CFM on a PBB-TE virtual switch

20 Enable CFM globally.

```
cfm enable
```

21 Create the CFM service for the PBB-TE virtual switch.

```
cfm service create static-pbt <static-pbt> name
<String[45]>
```

where

static-pbt <static-pbt> is the PBT tunnel service.

name <String[45]> is the CFM service name.

22 Enable the CFM PBB-TE service.

```
cfm service enable service <service>
```

where

service <service> is the service name.

To configure CFM on a PBB-TE tunnel

23 Enable CFM globally.

```
cfm enable
```

24 Create the CFM service for the PBB-TE tunnel.

```
cfm service create static-pbt <static-pbt> name
<String[45]> {next-mepid <NUMBER>]
```

where

static-pbt <static-pbt> is the PBT tunnel service.

name <String[45]> is the CFM service name.

next-mepid <NUMBER> initialize next-mepid

Copyright© 2012-2013 Ciena® Corporation

25 If desired, the CCM interval can be configured, but the interval must be the same on both sides of the tunnel endpoints. The default is 1 second.

Note: Setting the CCM Interval to 4 milliseconds is not recommended for a production environment, although it can be used for testing.

```
cfm service set service <service> ccm-interval <0 | 1 | 2
| 3 | 4 | 5 | 6 | 7 | off | 3.33ms | 10ms | 100ms | 1sec
| 10sec | 1min | 10min>
```

where

service <service>	is the service name.
ccm-interval <0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | off | 3.33ms | 10ms | 100ms | 1sec | 10sec | 1min | 10min>	is the CCM transmission interval index or time.

26 Enable the CFM PBB-TE service.

```
cfm service enable service <service>
```

where

service <service>	is the service name.

27 Verify that the MEP IDs at both ends of the tunnel are different:

```
cfm mep show
```

```
cfm remote-mep show
```

28 If the MEPS IDs are the same, change one of the MEP IDs:

```
cfm mep set service <service> port <Port> mepid <NUMBER:
1-8191>
```

where

service <service>	is the service name.
port <Port>	is the port.
mepid <NUMBER: 1-8191>	is the MEP ID.

To configure QoS on a PBB-TE tunnel

29 Set the port into advanced mode:

```
traffic-profiling set port <port> mode advanced
```

where

port <port>	is the port.

Copyright© 2012-2013 Ciena® Corporation

30 Set up ingress metering:

```
traffic-profiling standard-profile create port <port>
profile <NUMBER: 1-4096> cir <NUMBER> pir <NUMBER>
```

where

port <port>	is the port number.
profile <NUMBER: 1-4096>	is the traffic profile number.
cir <NUMBER>	is the committed information rate in Kbps of this traffic meter, rounded to fit hardware.
pir <NUMBER>	is the peak information rate in Kbps of this traffic meter, rounded to fit hardware

```
traffic-profiling set port <port> mode <none | advanced |
standard-dot1dpri | standard-ip-prec | standard-dscp |
standard-vlan | standard-vlan-dot1dpri | standard-vlan-
ip-prec | standard-vlan-dscp | hierarchical-port |
hierarchical-vlan>
```

where

port <port>	is the port.
mode <none | advanced | standard-dot1dpri | standard-ip-prec | standard-dscp | standard-vlan | standard-vlan-dot1dpri | standard-vlan-ip-prec | standard-vlan-dscp | hierarchical-port | hierarchical-vlan>	is the remarking mode.

Copyright© 2012-2013 Ciena® Corporation

31 Set CoS marking:

```
tunnel encap set static-pbt <static-pbt> encap-cos-policy
<fixed> encap-fixed-pcp <NUMBER: 0-7>
```

where

static-pbt <static-pbt>	is the backup tunnel name.
encap-cos-policy <fixed>	is the encapsulation CoS policy.
encap-fixed-pcp <NUMBER: 0-7>	is the encapsulation fixed PCP policy.

```
virtual-switch ethernet set vs <vs> {decap-cos-policy
<fixed | leave | tunnel-inherit | vc-inherit>} {decap-
fixed-dot1dpri <NUMBER: 0-7>} {description <String[128]>}
{encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-
inherit | phbg-inherit>} {encap-fixed-dot1dpri <NUMBER:
0-7>} {subscriber-dot1dpri-policy <leave | provider-
inherit>} {vc <Virtual Circuit Name>} [ip-iterface
<interface-name>]
```

where

vs <vs>	is the virtual switch name
encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit>	is the encapsulation CoS policy.
encap-fixed-dot1dpri <NUMBER: 0-7>	is the encapsulation fixed dot1dpri.

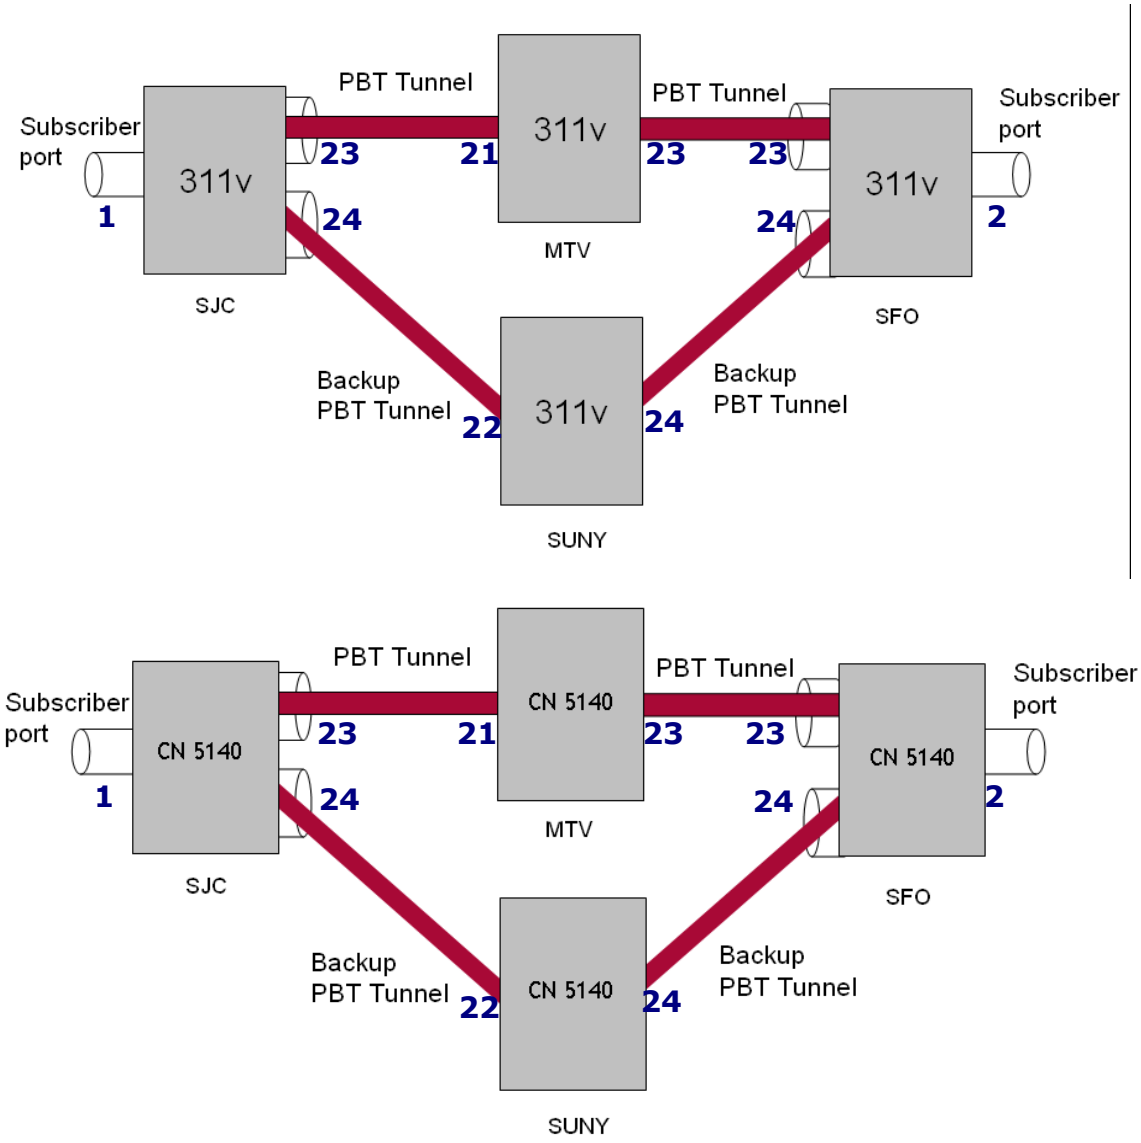
—end—

Example

The following configuration example displays the steps to configure a primary and a backup PBB-TE tunnel over the devices shown in Figure 10-7. It also employs CFM to monitor the connection. It assumes that PBB-TE mode is enabled and that a valid PBB-TE license key is installed.

Copyright© 2012-2013 Ciena® Corporation

Figure 10-7
PBB-TE Configuration Example



The following example configures device 1 SJC.

```
SJC> rstp disable -OR- SJC> mstp disable
SJC> pbt set bridge-mac 00:02:a1:30:7c:80
SJC> pbt reserve bvid 2000,3000
SJC> pbt remote-bridge create bridge-name SFO bridge-mac 00:02:a1:30:7c:c0
SJC> tunnel encap create static-pbt to_SFO b-vid 2000 dest-bridge-name SFO
port 23 tunnel-sync on
SJC> tunnel decap create static-pbt frm_SFO b-vid 2000 src-bridge-name SFO
port 23
SJC> tunnel pair create tnl-pair SJC encap-pbt to_SFO decap-pbt frm_SFO
SJC> tunnel encap create static-backup-pbt to_SFO b-vid 3000 dest-bridge-name
SFO port 24
```


Copyright© 2012-2013 Ciena® Corporation

```
SJC> tunnel decap create static-pbt bkp_SFO b-vid 3000 src-bridge-name SFO
port 24
SJC> tunnel pair create tnl-pair SJC-Backup encap-backup-pbt to_SFO decap-pbt
bkp_sfo
SJC> virtual-circuit pbt create static-vc pbt dest-bridge-name SFO egress-isid
1 ingress-isid 2 tunnel to_SFO
SJC> virtual-switch add reserved-vlan 1000
SJC> virtual-switch ethernet create vs pbt vc pbt
SJC> virtual-switch ethernet add vs pbt port 1 vlan 10
SJC> cfm service create static-pbt to_SFO name PBT level 7 next-mepid 100
SJC> cfm service set service PBT alarm-time 0
SJC> cfm service enable service PBT
SJC> cfm service create static-backup-pbt to_SFO name PBT_BKP level 7 next-
mepid 101
SJC> cfm service enable service PBT_BKP
```

The following example configures device 2 SFO.

```
SFO> rstp disable -OR- SFO> mstp disable
SFO> pbt set bridge-mac 00:02:a1:30:7c:c0
SFO> pbt reserve bvid 2000,3000
SFO> pbt remote-bridge create bridge-name SJC bridge-mac 00:02:a1:30:7c:80
SFO> tunnel encap create static-pbt to_SJC b-vid 2000 dest-bridge-name SJC
port 23 tunnel-sync on
SFO> tunnel decap create static-pbt to_SJC b-vid 2000 src-bridge-name SJC port
23
SFO> tunnel pair create tnl-pair SFO encap-pbt to_SJC decap-pbt to_SJC
SFO> tunnel encap create static-backup-pbt to_SJC b-vid 3000 dest-bridge-name
SJC port 24
SFO> tunnel decap create static-pbt to_SJC b-vid 3000 src-bridge-name SJC port
24
SFO> tunnel pair create tnl-pair SFO-Backup encap-backup-pbt to_SJC decap-pbt
to_SJC
SFO> virtual-circuit pbt create static-vc pbt dest-bridge-name SJC egress-isid
2 ingress 1 tunnel to_SJC
SFO> virtual-switch add reserved-vlan 1000
SFO> virtual-switch ethernet create vs pbt vc pbt
SFO> virtual-switch ethernet add vs pbt port 1 vlan 10
SFO> cfm service create static-pbt to_SJC name PBT level 7 next-mepid 200
SFO> cfm service set service PBT alarm-time 0
SFO> cfm service enable service PBT
SFO> cfm service create static-backup-pbt to_SJC name PBT_BKP level 7 next-
mepid 201
SFO> cfm service enable service PBT_BKP
```

The following example configures device 3 MTV.

```
MTV> rstp disable -OR- MTV> mstp disable
MTV> vlan create vlan 2000,3000
MTV> vlan add vlan 2000,3000 port 21,23
MTV> cfm mip create vlan 2000 port 21 level 6
MTV> cfm mip create vlan 2000 port 23 level 6
MTV> cfm mip create vlan 3000 port 21 level 6
MTV> cfm mip create vlan 3000 port 23 level 6
```

The following example configures device 4 Suny:

Copyright© 2012-2013 Ciena® Corporation

```
Suny> rstp disable -OR- Suny> mstp disable
Suny> vlan create vlan 2000,3000
Suny> vlan add vlan 2000,3000 port 22,26
Suny> cfm mip create vlan 2000 port 22 level 6
Suny> cfm mip create vlan 2000 port 24 level 6
Suny> cfm mip create vlan 3000 port 22 level 6
Suny> cfm mip create vlan 3000 port 24 level 6
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 10-6
Releasing reserved BVIDs

Release reserved BVIDs If you removed the PBB-TE configuration and want to reuse the BVID for a VLAN ID.

Step	Action
1	Release reserved BVIDs: `pbt release bvid <VLAN ID List>` where bvid <VLAN ID List> is the VLAN or VLANs to be released.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 10-7
Displaying PBB-TE information

Display PBB-TE information to view the PBB-TE global configuration.

Step	Action
1	Display PBB-TE global configuration information: `pbt show`

—end—

Example

The following example shows sample configuration for the pbt show command.

```
+---------------------- PBT Attributes -----------------------+
| PBT Bridge MAC                    | 00:03:18:cf:ea:80       |
| PBT Service E-type                | 0x88c8                  |
| PBT Tunnel E-type                 | 0x88a8                  |
| PBT S-tag Tpid                    | 0x8100                  |
| PBT Transit Tunnel E-type Remark  | On                      |
| PBT Tunnel Reversion              | Off                     |
| PBT Reversion Hold Time           | 3000       milli secs   |
| PBT Admin Mode                    | Non-native              |
| PBT Oper Mode                     | Non-native              |
+-----------------------------------+-------------------------+

+---------------- PBT Reserved B-VIDs -------------------+
| B-VID                        | In Use                  |
+------------------------------+-------------------------+
| No entries found                                       |
+------------------------------+-------------------------+

+---------------- PBT Remote Bridges --------------------+
| Bridge Name      | Index | Bridge MAC         | In use |
+------------------+-------+--------------------+--------+
| No entries found                                       |
+------------------+-------+--------------------+--------+
```

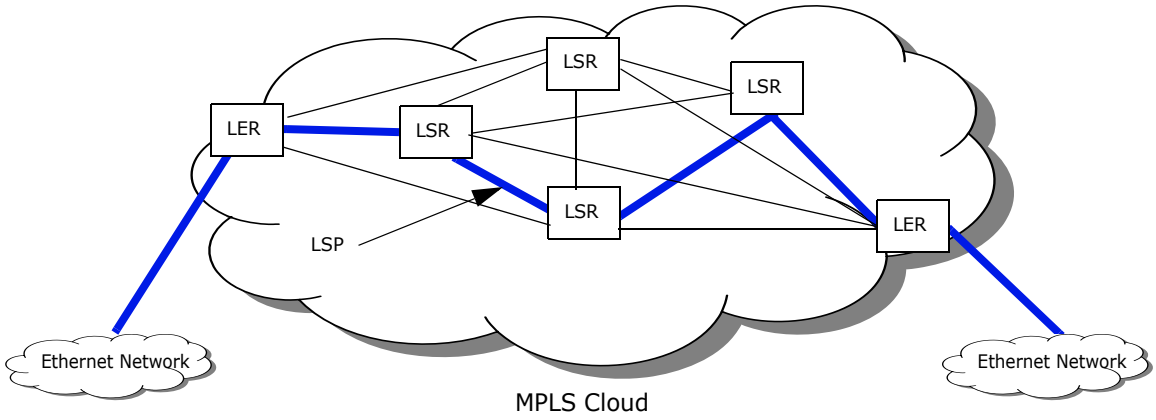

Copyright© 2012-2013 Ciena® Corporation

Copyright© 2012-2013 Ciena® Corporation

Multiprotocol Label Switching (MPLS) Layer 2 VPN configuration

An MPLS network (or cloud) is defined by a region of two MPLS enabled switches or routers called Label Edge Routers (LERs), one for ingress and one for egress. IP traffic enters an MPLS network through the ingress LER, which assigns labels based upon the forwarding equivalency class (FEC) it belongs to and initiates the RSVP-TE signaling to create the Label Switched Path (LSP). These labels are forwarded over the LSP to the Label Switch Routers (LSRs), where labels are switched to the next hop in the LSP. When the packets reach the egress LER, the MPLS labels are removed and the IP packets are delivered.

Figure 11-1
MPLS network

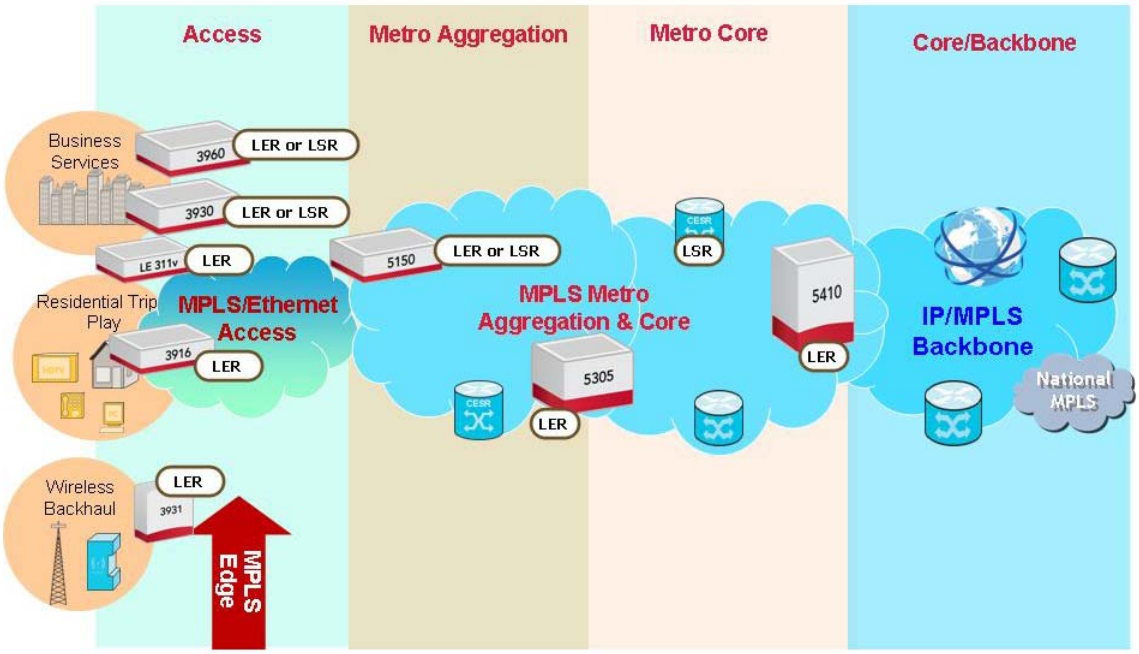


Transport of Ethernet/802.3 and VLAN tagged (802.1Q) traffic between multiple sites that belong to the same L2 broadcast domain or MPLS service occurs over MPLS PseudoWires (virtual circuits). MPLS virtual circuits can be configured statically or signaled using Label Distribution Protocol (LDP). The virtual circuits are tunneled and the tunnels can be configured statically or signaled using RSVP-TE.

Copyright© 2012-2013 Ciena® Corporation

The 3916, 3930, 3931, 3960, and 5150 switches are able to operate as an LER or LSR enabling MPLS based L2 VPNs to be implemented from the backbone to the access network.

Figure 11-2
MPLS Edge in the Access Network



Terms and definitions

This section provides a definition of terms used to describe MPLS networks.

Term	Definition
Attachment Circuit (AC)	Represents the client Layer 2 circuit on the UNI port of the PE router.
MTU-s	Multi-tenant Unit Switch
Provider Edge (PE) Router	Router responsible for encapsulating client data to be carried across the MPLS tunnels.
Provider Router	LSR that operates on the MPLS tunnel label.
PseudoWire (PW)	Virtual circuit that encapsulates the client payload and adds the PW header.

Copyright© 2012-2013 Ciena® Corporation

System requirements

To enable MPLS and related protocol commands, the MPLS feature license must be installed with the `software license install` command.

MPLS is supported on 3916, 3930, 3931, 3960, and 5150 switches on all ports as of the SAOS 6.9 release.

MPLS services

Ciena's implementation of MPLS is based on the Martini draft for L2VPN label exchange. This indicates that Ciena uses LDP for label exchange rather than Border Gateway Protocol (BGP). Additional information on label exchange can be found in "Signaling protocols" on page 11-11.

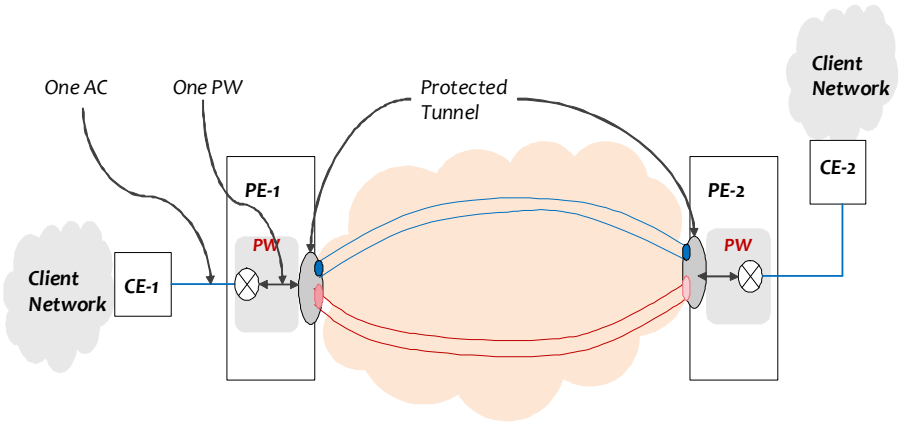
This implementation supports the simultaneous deployment of MPLS based L2 VPN services, including:

- Virtual Private Wire Services (VPWS).
- Virtual Private Local Area Network Services (VPLS).
- Hierarchical VPLS (H-VPLS).

VPWS

VPWS provides point-to-point connectivity between two remote Local Area Networks (LANs). With this type of connectivity, MPLS provides the equivalent of an E-LINE service (EPL or EVPL). Traffic is carried from a single AC to exactly one PW, and MAC addresses are not learned for the Ethernet AC. Figure 11-3 provides an example of a VPWS with a one-to-one connection between CE1 and PE1.

Figure 11-3
VPWS connectivity

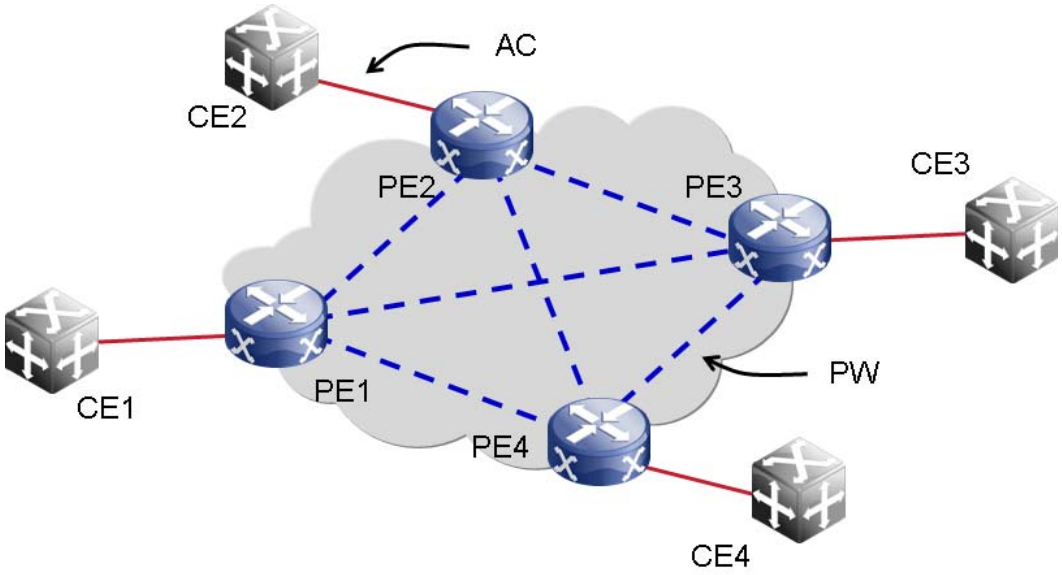


Copyright© 2012-2013 Ciena® Corporation

VPLS

VPLS provides point-to-multipoint inter-LAN connectivity. With VPLS, PEs are connected to each other with a full mesh of virtual circuits for each VPLS instance as shown in Figure 11-4.

Figure 11-4
VPLS full mesh



Each PE provides one or more connections to the customer edge (CE) devices via an AC to the PE MAC/VLAN bridge function. The PE devices are connected to each other using virtual circuits. The PE device's bridge function attaches to an emulated LAN, defined by a virtual switch, consisting of "VPLS Forwarders" which connect to virtual circuits. This set of VPLS Forwarders connected by virtual circuits represents a VPLS instance.

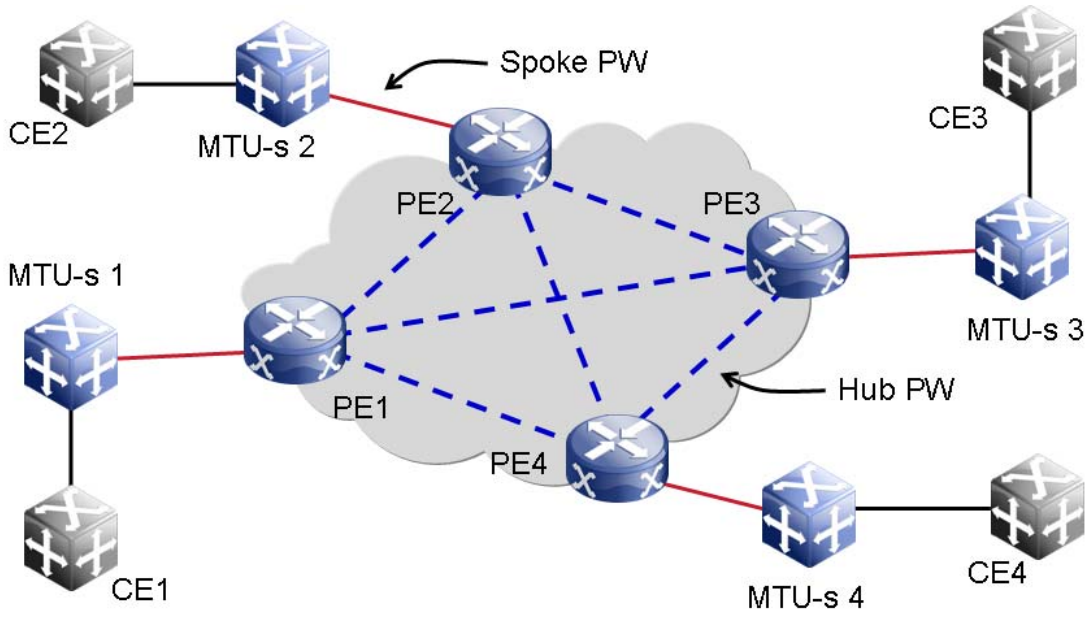
The bridge function learns MAC addresses, associates them with virtual circuits, and ages them out in a standard manner. When a broadcast/multicast/unknown frame arrives at the bridge function, the frame is forwarded over all the virtual circuits attached to the VPLS forwarder. When a frame arrives on a virtual circuit, the bridge function performs normal Source Address (SA) MAC learning to associate the MAC address with the virtual circuit. In this way, the PE device emulates the behavior of a normal LAN. The CE devices appear to be connected to a LAN even though the underlying infrastructure is an MPLS network.

Copyright© 2012-2013 Ciena® Corporation

H-VPLS

The deployment of large-scale VPLS networks where each PE is connected to all other PEs using a full mesh of Virtual Circuits does not allow complete scalability. As a result, a Hierarchical VPLS (H-VPLS) model is used to allow spoke connections to the VPLS core. As shown in Figure 11-5, a device provides the functionality to interface with the VPLS core by functioning as an MTU-s, which is connected as a spoke in the VPLS core using a virtual circuit.

Figure 11-5
H-VPLS with MTU-s spoke



In this H-VPLS model, there is only one logical connection (virtual circuit) from an MTU-s to the PE for a given VPLS instance. Each VPLS instance supported by an MTU-s has a virtual switch defined as a virtual Layer 2 switch.

The MTU-s must map incoming frames on its bridge module into VPLS instances. This can be based on one of three methods:

- Physical ports.
- LAN tag of the ingress frame.
- MPLS label (VC label) of the ingress frame.

Copyright© 2012-2013 Ciena® Corporation

VPLS membership and MAC learning

The virtual switch provides a mapping between customer-facing ACs and network-facing virtual circuits. A virtual switch may be dual-homed such that there are multiple virtual circuits connecting the VPLS instance to the VPLS hub. An AC can be defined on a per-port or per-port/per-VLAN basis depending upon the classifications of the sub-port interface. MAC addresses to be included in forwarding for the VPLS instance are learned in one of two ways:

- Unqualified Learning. In unqualified learning, all traffic from a specific bridge port is assigned to a single VPLS instance (per-port AC) and shares a single broadcast domain. MAC addresses need to be unique and non-overlapping among customer VLANs or else they cannot be differentiated within the VPLS instance.

  An example of this is port-based VPLS service for a given customer. In this case, all the traffic that comes in on a physical port (which may include multiple customer VLANs) is mapped to a single VPLS instance.

- Qualified Learning. In qualified learning, each customer VLAN is assigned to its own VPLS instance, which means that each customer VLAN has its own broadcast domain and MAC address space. Unlike unqualified learning, the MAC addresses among customer VLANs can overlap with each other. They are handled correctly since each customer VLAN has its own Forwarding Information Base (FIB). When a VPLS instance is defined per-port/per-VLAN, the customer VLAN must be the same on each bridge port that is joining the VS. An example of qualified learning is per-port-per-VLAN VPLS service.

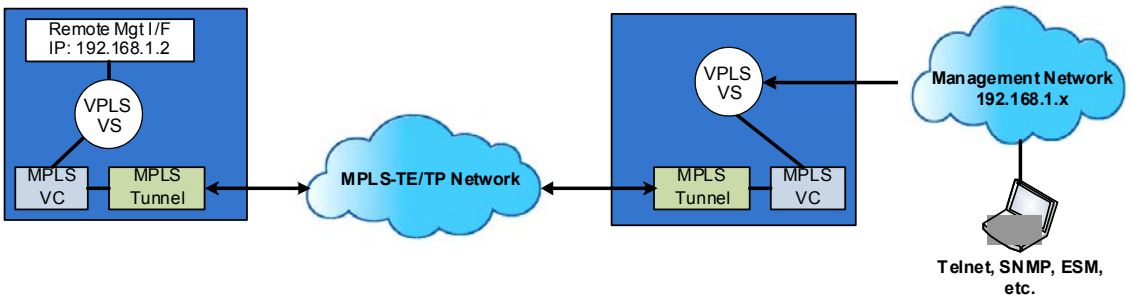
Remote Management for MPLS

You can access an IPv4 Remote Management Interface over an MPLS VPLS mode Virtual Switch (VS) as an alternative to the management VLAN.

Before configuring management over MPLS, you must first allocate resources to the transport-oam feature on 3916/30/31 platforms. For more information, see “Configuring transport OAM resources”.

Copyright© 2012-2013 Ciena® Corporation

Figure 11-6
Remote Management Interface over an MPLS VPLS VS



In order to carry remote interface traffic over an MPLS tunnel, the remote-interface is associated with a Virtual Switch (VS). One or more MPLS VCs can belong to the management VS, allowing the management traffic to travel over the MPLS tunnel(s).

This capability is supported on the following platforms:

- 3916
- 3930
- 3931
- 5150

In order to carry remote interface traffic over an MPLS tunnel, the remote-interface can be associated with a VS. This implies that management access to the switch can now be gained from any of the members of this VS, including AC members. Thus, if customer ACs exist on the VS that is associated with the Remote Interface, the customers may be able to obtain management access to the node. In order to prevent this, it is recommended that the service provider create an MPLS VS specifically for use for management, and only add any AC members to the VS in the following scenario: on a boundary node where the network transitions from VLAN based management, to in-band MPLS based management, a port where the VLAN based management traffic arrives may be added to the MPLS VS as an EVPL AC.

Interactions with PBT mgmt-vc

Legacy PBT in-band management supports creation of a mgmt-vc to provide management access over PBT tunnels. Remote interface management access over PBT is mutually exclusive with Remote interface management access over MPLS. The network operator cannot configure a PBT mgmt-vc if an MPLS VS is currently set on the remote interface, and vice-versa.

Note: MPLS virtual switch in-band management does not support IPv6.

Copyright© 2012-2013 Ciena® Corporation

In order to change from management over PBT to management over MPLS, delete the mgmt-vc before issuing the "interface remote set vs" command to set the MPLS VS.

Alternately, to replace an MPLS VS already associated with the remote interface with management over PBT, the management must be moved to a VLAN using the "interface remote set vlan <xxx> command. Only then can the PBT mgmt-vc command be invoked.

> ***Note:*** In native PBT in-band management, you can use the interface remote set <vlan|vs> command to switch between a VLAN, MPLS virtual switch and PBT virtual switch.

Changing the management VS

You can create several MPLS VSs, although only one VS can be associated with the remote interface at any point in time. Even if the remote interface is already associated with an MPLS VS, you can create a new MPLS VS and change to using this VS for management by associating the new VS with the remote interface as follows:

```
>interface remote set vs <new_vs_name>
```

At this point remote interface connectivity is moved to the new VS, and it becomes the “management VS”.

You can also replace the “management VS” with a “management VLAN” (or vice versa) at any time, using the “interface remote set” command.

For example:

```
>interface remote set vs <mpls-vs-name>
```
(remote mgmt is only through the VS)

```
>interface remote set vlan 127
```
(remote mgmt is now only through the VLAN)

```
interface remote set vs <mpls-vs-name>
```
(remote mgmt is now again only through the VS)

The old “management” VS (or VLAN) continues to exist and its membership is not changed, but it no longer provides connectivity to the remote interface.

Routing protocols

MPLS network topology and routing information is monitored and maintained by the following routing protocols:

- Open Shortest Path First (OSPF).
- Open Systems Interconnect (OSI) Intermediate System to Intermediate System (IS-IS) Intra-domain Routing Protocol.

Copyright© 2012-2013 Ciena® Corporation

These routing protocols discover neighbors in the MPLS network and maintain the optimum route to them.

OSPF

Open Shortest Path First (OSPF) is used to create an IP routing table for building dynamic MPLS tunnels. One of several Interior Gateway Protocols (IGPs) designed to support routing in an IP network within a single autonomous system (AS), OSPF is a link-state protocol that uses configurable metrics to associate a cost with a device. These metrics allow network administrators to manage their network based on the speed, reliability, and delay of the network.

The OSPF protocol is a link-state routing protocol, which means that the routers exchange topology information with their nearest neighbors. The topology information, in the form of a link-state advertisement (LSA), is flooded throughout the AS, so that every router within the system has a complete representation of the topology. Devices build a Link State Database (LSDB) based on this information. Each Area Border Router (ABR) has one LSDB for each area to which it is connected. This information is then used to calculate the shortest end-to-end paths through the system. This is accomplished by means of Djikstra's Shortest Path First (SPF) algorithm. The SPF tree, also known as the best path tree, is then submitted to the routing table as OSPF routes. When a network topology change occurs, a recalculation of the shortest path tree must be performed. The network has converged when all routers have recalculated their routing tables as a result of a change in the topology.

OSPF neighbors are any two routers that have an interface to the same network. When an OSPF device first joins a network, it uses the Hello Protocol to discover its neighbors. Neighbors may form adjacencies for the purpose of exchanging routing information. Not all neighbor pairs can become adjacent. Adjacencies are formed by synchronizing the neighbors' topology databases through the database exchange process. Two devices are said to be fully adjacent when they have synchronized their databases. Routing information is exchanged between the adjacent routers only, thereby conserving bandwidth. Also, an authentication mechanism prevents unauthorized neighbors from establishing adjacencies.

The multi-level hierarchy (two-level for OSPF) called “area routing” allows the information about the topology within a defined area of the AS to be hidden from routers outside this area, enabling an additional level of routing protection and a reduction in routing protocol traffic. The authentication of all protocol exchanges prevents unauthorized routers from joining the AS.

Copyright© 2012-2013 Ciena® Corporation

IS-IS

Similar to OSPF, IS-IS is a link-state routing protocol in the group of IGPs and calculates the shortest end-to-end paths through the system with Djikstra's Shortest Path First (SPF) algorithm. As defined in RFC1195, IS-IS supports routing within IP only, OSI only, and dual (IP with OSI) domains and supports two router levels, Level 1 and Level 2. With IS-IS, a routing domain can be partitioned into multiple Level 1 router areas with a Level 2 interconnection router. Level 1 routers within the same area exchange information, but for destinations outside of the area, Level 1 routers forward to the nearest Level 2 or dual (Level 1 and Level 2) router.

On a broadcast, multi-access network, an intermediate system is elected as a designated router. An intermediate system supports one designated router on a broadcast link for Level 1 and one for Level 2. The designated router on a broadcast link represents a virtual node, called a Pseudo Node.

The Pseudo Node creates and updates link state packets to report link state information to all systems on the broadcast sub-network. Also, it floods link state packets on the network and periodically sends Complete Sequence Numbers Packets (CSNPs) to synchronize link state databases at each intermediate system on the network. All intermediate systems, including the designated router, form adjacencies with the Pseudo Node using IS-IS Hello Packets (IIH). To acknowledge receipt or to request link state information, intermediate systems send Partial Sequence Number Packets (PSNP).

Intermediate systems are identified by a unique Network Entity Title (NET) composed of the following:

- Area identifier. Variable dotted hexadecimal value from 2 bytes to 13 bytes with the first byte value containing the Address Format Identifier (AFI).
- System identifier. A 6 byte hex value derived from the loopback IP address.

The link state packet includes the NET and may optionally include a Dynamic Hostname TLV containing the symbolic name of the router, such as a Fully Qualified Domain Name, sending the link state packets.

In this release, the system software supports Level 1 routing capability for IPv4 protocol and 1 area with 3 area identifiers per node. IS-IS can be used to replace OSPF functionality or it can be implemented simultaneously with OSPF.

Copyright© 2012-2013 Ciena® Corporation

Signaling protocols

MPLS networks are created and monitored with the signaling of Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for tunnels, and Label Distribution Protocol (LDP) for virtual circuits.

Tunnels and virtual circuits use label exchange following Martini encapsulation without Control Word as defined in RFC4447. Martini encapsulation consists of adding an Ethernet transport header to the beginning of an incoming packet that includes the following:

- Destination MAC address (DA) of the Provider Edge (PE) device is contained in the Ethernet transport header when directly connected to a device.
- Source MAC address of the Multi-tenant Unit (MTU) or PE and the provider VLAN.
- MPLS Ethernet-type and both the tunnel and virtual circuit MPLS labels.

RSVP-TE

RSVP-TE is a protocol used to establish label switched paths (LSPs) in MPLS networks. These LSPs allow the traffic trunks, a set of flows aggregated by their service class on one or a set of LSPs, to be carried through the network. One LSP can carry several traffic trunks. The traffic that flows along an LSP is defined at the ingress node. When labels are associated with traffic flows, routers can identify the appropriate reservation state for a packet. These defined paths can be treated as tunnels, and are described as LSP tunnels.

Tunnels

LSP tunnels can be used to implement various policies to optimize network performance. For example, policies can be established to reroute traffic when there is a node failure, or when a path is overloaded with packets. RSVP-TE tunnels can be established along a Traffic Engineered path identified using Constrained Shortest Path First (CSPF is an algorithm used to resolve Quality of Service routing queries) or tunnels can be established by following a user-configured explicit path to the destination device. A Tunnel is defined as a unidirectional LSP path between two devices.

Failover

MPLS tunnels can be protected by configuring backup tunnels at the head-end LER. The backup tunnel provides protection end-to-end for the primary tunnel via a set of LSR hops that are primary-path-negated. This ensures that failure on primary path does not translate to simultaneous failure on the backup path.

Copyright© 2012-2013 Ciena® Corporation

LDP

The principal role of LDP is to establish and maintain virtual circuits based upon the agreement of the meaning of the label used to forward traffic.

Because LDP is a peer-to-peer protocol based on the establishment and maintenance of TCP sessions, the following natural benefits exist:

- LDP messages are reliably delivered by the underlying TCP, and State information associated with explicitly routed LSPs does not require periodic refresh.
- LDP messages are flow controlled (throttled) through TCP.

The categories of LDP messages include:

- Discovery messages announce and maintain the LSR presence in the network.
- Session messages create, maintain and terminate sessions between two LDP peers that use LDP to exchange label mapping information.
- Advertisement messages create, change and delete label mapping for Forwarding Equivalence Class (FEC), a set of packets with similar or identical characteristics.
- Notification messages provide advisory information and signal error information.

Connectivity Fault Management over MPLS

MPLS virtual switches are monitored through the use of IEEE 802.1ag Connectivity Fault Management (CFM) Continuity Check Messages (CCM).

Note: If a CFM service is configured on a VS, the VS cannot be deleted unless the CFM service is deleted or changed to use a different VS.

Bidirectional Forwarding Detection (BFD)

MPLS-BFD implements a single-session BFD over an associated bidirectional LSP, constructed from a pair of unidirectional LSPs, one in each direction. The two LSPs operate virtually as one single entity for the purpose of protection and switching.

For additional details about BFD, refer to Fault and Performance (009-3203-009).

Complementary protocols

Complementary protocols include:

- LSP ping
- LSP traceroute

Copyright© 2012-2013 Ciena® Corporation

LSP ping

LSP ping provides the ability to verify connectivity and detect faults of RSVP-TE tunnels through the exchange of standard Echo Request and Echo Reply messages.

LSP traceroute

LSP traceroute provides the ability to verify the path and isolate faults of RSVP-TE tunnels through the exchange of standard Echo Request (with increment TTL) and Echo Reply (with downstream mapping from transit nodes).

Benefits

With MPLS and related protocols, service providers can provide customers with the perspective of a direct connection to a private line or LAN between their sites with enhanced performance, topology discovery to ensure the most efficient routes, and interoperability with popular MPLS switching platforms.

Vendor interoperability

This MPLS implementation supports configurable label ranges for both static and dynamic MPLS, and provides interoperability with other vendor devices.

Configuring L2 VPNs

The general steps for configuring MPLS are:

1 Configure the IP interfaces.
2 Configure the OSPF or IS-IS routing protocol(s).
3 Configure the RSVP-TE protocol.
4 Configure the RSVP-TE paths.
5 Configure dynamic or static tunnels.
6 Configure LDP.
7 Configure the virtual circuits.

For each MPLS deployment scenario, an IP interface must be configured for an existing VLAN mapped over a physical port for handling MPLS control protocol traffic. The 3916, 3930, and 3931 switches support 64 IP interfaces. The 3960 and 5150 switches support 256 IP interfaces.

Note 1: A different VLAN should be used for each IP Interface and physical port combination participating in the VPLS to prevents the creation of flood domains across Layer 3 segments.

Note 2: The loopback address is typically used as the router identifier for routing protocols, LDP identifier, MPLS tunnel peer address, and more.

Copyright© 2012-2013 Ciena® Corporation

Virtual circuits

A virtual circuit is a bidirectional connection between endpoints which can multiplex/de-multiplex traffic over tunnels. Multiple virtual circuits between two endpoints may use the same tunnels. For virtual circuits associated with an MTU-s, a secondary virtual circuit can be configured to support dual-homed protection.

Note: The 3960, 3916, and 3931 switches support up to 125 virtual circuits and 125 virtual switches. The 3930 and 5150 switches support up to 250 virtual circuits and 250 virtual switches.

Primary virtual circuit attributes

Primary dynamic and static MPLS virtual circuits support the attributes as listed in Table 11-1.

Table 11-1
Primary virtual circuit attributes

Attribute	Description
Name	Unique 31 character name.
PW-ID	PseudoWire identifier. Ranging from 1-4294967295. This identifier must be the same for all virtual circuits associated with a role when operating with other vendor equipment.
Peer	Peer IP address.
Ingress label	Ingress label for the virtual circuit. Only applies to static virtual circuits.
Egress label	Egress label for the virtual circuit. Only applies to static virtual circuits.
Tunnel	Name of the tunnel.
PW type	PseudoWire type. For static virtual circuits, the types are raw or tagged. Default is raw. For dynamic virtual circuits, the types are eth-raw or eth-tagged.
PW mode	PseudoWire mode: • Mesh. A virtual circuit between two PEs. • Spoke. A virtual circuit emanating towards the MTU-s.
MTU	Maximum transmission unit 1500-9128. Default is 1500. This parameter is exchanged during dynamic establishment of a PW and must match in order for the PW to become operational between two PEs. ***Note:*** Mismatch of the MTU is one of the most common reasons for PW to fail to become operational.

Copyright© 2012-2013 Ciena® Corporation

Table 11-1
Primary virtual circuit attributes

Attribute	Description
Status TLV	On or off. LDP notifications carry status. Not supported.
Service Delimiter VID	Supports per VLAN TPID stamping for the specified VLAN ID (1-4094). Used with the Service Delimiter TPID.
Service Delimiter TPID	Selects the TPID value to stamp on egress frames for the VLAN ID specified with the Service Delimiter VID: • 8100 - 802.1Q Ethertype • 9100 - Q-in-Q Ethertype • 88a8 - Provider Bridge Ethertype

Comparing raw and tagged PW type for virtual circuits

The term Attachment Circuit (AC) is used to describe the customer facing port of an Ethernet L2VPN service. The 2 types of ACs available are EPL AC and EVPL AC. EPL AC describes when an entire port is added to an L2VPN service. EVPL AC describes when a single VLAN is added to an L2VPN service.

MPLS forwarding behavior can vary depending on the type of AC/VC combination configured. There are 4 possible combinations:

- EPL AC with raw virtual circuit
- EVPL AC with raw virtual circuit
- EPL AC with tagged virtual circuit
- EVPL AC with tagged virtual circuit

An EPL AC indicates that the outer VLAN tag of the frame, if present, is not service delimiting and therefore is not meaningful to the PE. An EPL AC always assumes the outer tag of the frame is a C-Tag.

An EVPL AC indicates that the outer VLAN tag of the frame is service delimiting and should be used to identify the traffic. An EVPL AC always assumes the outer tag of the frame is an S-Tag.

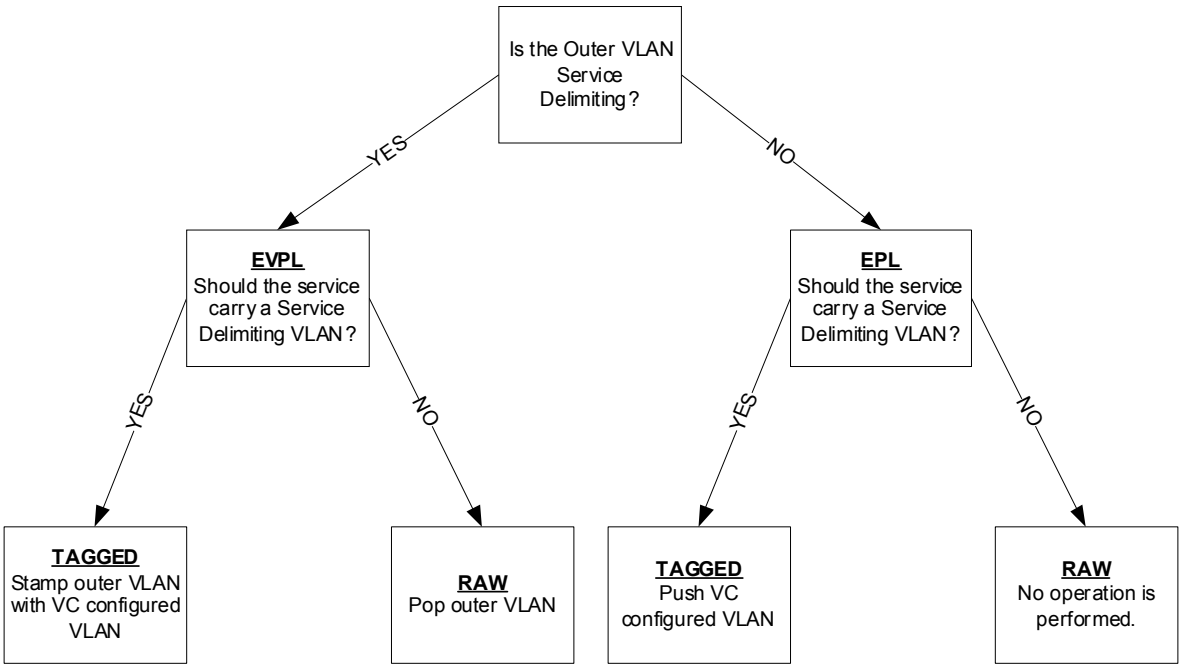
A raw PW type virtual circuit never carries a service delimiting tag. There are only 2 possible options for this type: ignore the tag or pop it. In the case of an EPL AC the tag is assumed to be a customer tag and it is ignored. In the case of an EVPL AC the tag is assumed to be a service provider tag and it is popped.

Copyright© 2012-2013 Ciena® Corporation

A tagged PW virtual circuit always carries a service delimiting tag. There are two possible options for this type: stamp or push. In the case of an EPL AC the outer tag is assumed to be a C-Tag and an additional tag is pushed on the frame. In the case of an EVPL AC the outer tag is assumed to be a service provider tag and it is stamped.

Figure 11-7 illustrates process for determining the PW type.

Figure 11-7
Ingress Operation Decision Tree



Note: If identical settings are configured on either end of the MPLS tunnel the original frame will be preserved through the service. The operation as the customer frame egresses the service is contrary to the ingress operation. However, the outer tag can also be permanently altered by changing the AC/virtual switch combination on the egress side.

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-1
Building an OSPF configuration

Build an OSPF configuration according to network requirements. Building an OSPF configuraiton comprises:

- create
- enable
- attach

In dynamic MPLS deployments, OSPF or IS-IS is required. In static MPLS deployments, OSPF is optional. The system supports up to 1000 OSPF routes.

> ***Note:*** You can configure both IS-IS and OSPF. However, to simplify implementation, using one of them is recommended.

The system supports a single OSPF instance with unique 31 character name, which is required to enable OSPF. By default, this OSPF instance is disabled.

The default configuration for an OSPF instance is not compatible with RFC 1583 when calculating OSPF based routes, and multiple intra AS routes are maintained. When AS boundary routers belong to different OSPF areas, and advertise the same external destination, enable support for RFC 1583 compatibility. Also, the default router ID when the instance is created is the Loopback address of the system and the SPF delay interval is 5000 milliseconds.

Step	Action

To create an OSPF instance

1 Create an OSPF instance:

```
ospf instance create ospf-instance <instance-name>
```

where

ospf-instance <instance-name> is the OSPF instance name.

Copyright© 2012-2013 Ciena® Corporation

To enable an OSPF instance

2 Enable an OSPF instance:

```
ospf instance enable ospf-instance <instance-name>
```

where

ospf-instance <instance-name>	is the OSPF instance name.

To enable OSPF for an IP interface

3 Enable OSPF on an IP interface:

```
ospf interface enable ip-interface <ip-interface>
```

where

ip-interface <ip-interface>	is the IP interface name.

To attach an OSPF instance to an IP interface

4 Attach an OSPF instance to an IP interface:

```
ospf interface attach ip-interface <ip-interface> ospf-
instance <instance-name>
```

where

ip-interface <ip-interface>	is the IP interface name.
ospf-instance <instance-name>	is the OSPF instance name.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-2
Configuring an OSPF area

An OSPF area is defined by an IP address and with the default type of normal. By default, OSPF has the single area 0.0.0.0.

At factory settings area 0.0.0.0 does not display in CLI. It is created when you first attach an interface to area 0.0.0.0. Once area 0 is created it cannot be deleted.

You can configure OSPF to use this default area, or you can configure an additional non-backbone area enabling the system to perform the functions of an Area Border Router (ABR).

Note: With type normal, the continuous backbone area called area 0.0.0.0 is directly connected to every other area and is used for inter-area routing.

Step	Action

To create an OSPF area

1 Create an OSPF area:

```
ospf area create area-ip <ipv4 address> ospf-instance
<instance-name> type normal
```

where

area-ip <ipv4 address>	is the IP address of the new non-backbone area.
ospf-instance <instance-name>	is the OSPF instance.
type	is the type of area.

To display the OSPF area or statistics tables

2 To display the OSPF area or statistics tables:

```
ospf area show ospf-instance <instance-name> [area <area-
ip>][statistics]
```

where

area-ip <ip-address>	is the IP address of the area.
ospf-instance <instance-name>	is the OSPF instance.

Copyright© 2012-2013 Ciena® Corporation

To delete an OSPF area

3 Remove an OSPF area:

```
ospf area delete area-ip <IPv4 address> ospf-instance
<instance-name>
```

where

area-ip <IPv4 address>	is the area IP address.
ospf-instance <instance-name>	is the OSPF instance.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-3
Displaying OSPF information

Display OSPF information when troubleshooting. You can display the configuration and status of an OSPF instance as well as transmitted, received, dropped, and various LSA statistics.

You can:

- display the OSPF instance configuration or statistics
- clear the OSPF instance statistics
- display the OSPF area or statistics
- display a list of neighbors for the OSPF instance
- display a specific neighbor for the OSPF instance
- display the database for the OSPF instance

Step	Action

To display the OSPF instance configuration or statistics

1 Display the OSPF instance configuration or statistics:

```
ospf instance show ospf-instance <instance-name>
[statistics]
```

where

ospf-instance <instance-name> is the OSPF instance name.

To clear the OSPF instance statistics

2 Clear the OSPF instance statistics:

```
ospf instance clear ospf-instance <instance-name>
statistics
```

where

ospf-instance <instance-name> is the OSPF instance name.

Copyright© 2012-2013 Ciena® Corporation

To display the OSPF area or statistics

3 Display the OSPF area or statistics:

```
ospf area show ospf-instance <instance-name> [area
<areaid>][statistics]
```

where

ospf-instance <instance-name>	is the OSPF instance name.
area <areaid>	is the area IP address.
statistics	displays statistics.

To display a list of neighbors for the OSPF instance

4 Display a list of neighbors for the OSPF instance:

```
ospf neighbor show ospf-instance <instance-name>
neighbor-ip <ip-address> [detail]
```

where

ospf-instance <instance-name>	is the OSPF instance name.
neighbor-ip <ip-address>	is the neighbor IP.
detail	details OSPF neighbor information.

To display a specific neighbor for the OSPF instance

5 Display a specific neighbor for the OSPF instance:

```
ospf neighbor show ospf-instance <instance-name>
neighbor-ip <ip-address>
```

where

ospf-instance <instance-name>	is the OSPF instance name.
neighbor-ip <ip-address>	is the neighbor IP.

To display the database for the OSPF instance

6 Display the database for the OSPF instance:

```
ospf database show ospf-instance <instance-name> [router
| nssa-external | self-originate | network | asbr-
summary]
```

where

ospf-instance <instance-name>	is the OSPF instance name.
router	displays OSPF router lsa
network	displays OSPF network lsa

Copyright© 2012-2013 Ciena® Corporation

where

nssa-external	displays OSPF nssa-external lsa
asbr-summary	displays OSPF asbr-summary lsa
self-originate	displays OSPF self originate lsa

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-4
Modifying OSPF attributes for an IP interface

Modify OSPF attributes for an IP interface.

Note: In order for higher level protocols, such as RSVP-TE and LDP, to use the route information that OSPF provides, a loopback interface is required.

Step	Action

1 Modify OSPF attributes for an IP interface:

```
ospf interface set ip-interface <ip-interface> delay-
interval <SECONDS: 1-2147483647> hello-interval <SECONDS:
1-49999> poll-interval <SECONDS: 1-2147483647> priority
<0-255> retransmit-delay <SECONDS: 1-3000> transmit-delay
<MILLISECONDS: 1-429496799> cost-metric <NUMBER:
0..65535> authentication-type <none | text | md5>
password <Password String> [password-secret <String[62]>]
```

where

ip-interface <ip-interface>	is the IP interface name.
delay-interval <SECONDS: 1-2147483647>	is the delay interval of the interface in seconds.Default is 40 seconds.
hello-interval <SECONDS: 1-49999>	is the hello interval of the interface in seconds. Default is 10 seconds.
poll-interval <SECONDS: 1-2147483647>	is the poll interval for the interface in seconds. Default is 120 seconds.
priority <0-255>	is the priority used for the interface. Default is 1.
retransmit-delay <SECONDS: 1-3000>	is the retransmit delay interval for the interface in seconds. Default is 5 seconds.
transmit-delay <MILLISECONDS: 1-429496799>	is the transmit delay interval for the interface in msec. Default is 100 milliseconds.
cost-metric	is the cost the interface. Default is 30.

Copyright© 2012-2013 Ciena® Corporation

where

authentication-type <none | md5 | text>	is the authentication type. Default is none. Optionally, you can set the authentication type to use text or MD5 authentication.
password <Password String>	is the authentication password (1..8 characters). Default is blank. Required when the authentication type is set to text.
[password-secret <String[62]>]	is the encoded password. Default is blank. Required when the authentication type is set to md5.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-5
Clearing OSPF authentication for an OSPF IP interface

Clear OSPF authentication for an OSPF IP interface.

Step	Action

1 Clear OSPF authentication for an OSPF IP interface:

```
ospf interface set ip-interface <ip-interface>
authentication-type <none | md5 | text>
```

where

ip-interface <ip-interface>	is the IP interface name.
authentication-type <none | md5 | text>	is the authentication type. Default is none. Optionally, you can set the authentication type to use text or MD5 authentication.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-6
Displaying OSPF interface information

Display OSPF interface information. You can display:

- OSPF interface information
- OSPF IP interface information

Step	Action

To display OSPF interface information

1 Display OSPF interface information:

```
ospf interface show [detail]
```

where

[detail]	details OSPF interface information.

To display OSPF IP interface information

2 Display OSPF IP interface information:

```
ospf interface show ip-interface <ip-interface>
statistics
```

where

ip-interface <ip-interface>	is the IP interface name.
statistics	shows OSPF interface statistics.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-7
Creating an IS-IS instance with a single area

In dynamic MPLS deployments, OSPF or IS-IS are required. In static MPLS deployments, IS-IS is optional. The system supports up to 400 IS-IS routes.

Note: You can configure both IS-IS and OSPF. However, to simplify implementation, using one of them is recommended.

You can configure a single IS-IS instance with a unique 31 character string name with up to 3 areas for a Level 1 router. By default, the IS-IS instance is enabled upon creation. You can disable IS-IS protocol for a specific IS-IS instance and then enable it.

Step	Action

1 Create an IS-IS instance with a single area:

```
isis instance create isis-instance <isis-instance>
[logical-id <NUMBER>] area <ISIS area> [level <L1>]
```

where

isis-instance <isis-instance>	is the isis instance name.
logical-id <NUMBER>]	is the creation index used in configuration.
area <ISIS area>	is the area identifier, which has a variable length in Hex (AFI.xxxx.xxxx.....).
level <L1>	is the routing level.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-8
Deleting an IS-IS instance

Delete an IS-IS instance.

Step	Action
1	Delete an IS-IS instance: `isis instance delete isis-instance <isis-instance>` where isis-instance <isis-instance> is the isis instance name.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-9
Configuring IS-IS areas

You can

- add an IS-IS area
- remove an IS-IS area

Step	Action

To add an area

1 Add an IS-IS area:

```
isis instance add isis-instance <isis-instance> area
<ISIS area>
```

where

isis-instance <isis-instance>	is the isis instance name.
area <ISIS area>	is the area identifier, which has a variable length in Hex (AFI.xxxx.xxxx.....).

To remove an area

2 Remove an IS-IS area:

```
isis instance remove isis-instance <isis-instance> area
<ISIS area>
```

where

isis-instance <isis-instance>	is the isis instance name.
area <ISIS area>	is the area identifier, which has a variable length in Hex (AFI.xxxx.xxxx.....).

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-10
Enabling or disabling an IS-IS instance

Enable an IS-IS instance.

Disable an IS-IS instance.

Step	Action

To enable an IS-IS instance

1 Enable an IS-IS instance:

```
isis instance enable isis-instance <isis-instance>
```

where

isis-instance <isis-instance> is the isis instance name.

To disable an IS-IS instance

2 Disable an IS-IS instance:

```
isis instance disable isis-instance <isis-instance>
```

where

isis-instance <isis-instance> is the isis instance name.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-11
Displaying IS-IS instance protocol information

Display IS-IS instance protocol information.

Step	Action

1 Display IS-IS instance protocol information:

```
isis instance show isis-instance <isis-instance>
{attached-interfaces | database | hostname | is-neighbors
{summary | details} | neighbors {summary | details} |
statistics}
```

where

isis-instance <isis-instance>	is the isis instance name.
attached-interfaces	is the ISIS instance with its attachments.
database	is the protocol database details.
hostname	is the known hostnames of IS neighbors.
is-neighbors {summary | details}	is IS neighbors.
neighbors {summary | details}	is the IP address of neighbors.
statistics	is protocol statistics.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-12
Clearing IS-IS instance statistics

Clear IS-IS instance statistics.

Step	Action
1	Clear IS-IS instance statistics: `isis instance clear isis-instance <isis-instance> statistics` where isis-instance <isis-instance> is the isis instance name.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-13
Configuring IS-IS

IS-IS configuration comprises:

- "Link state packet generation"
- "Shortest Path First claculations"
- "Area authentication"
- "Interface authentication"

Link state packet generation

You can control the maximum lifetime of link state packets per instance as well as the refresh interval time between generation. The default maximum lifetime is 1200 seconds, and the default refresh interval is 900 seconds.

Shortest Path First claculations

You can control when updates to the link state database trigger Shortest Path First (SPF) routing calculations to occur with the following attributes:

- Maximum delay. Determines how long to wait after a database updated before updating SPF routing calculations. Default is 5000 milliseconds. When set to 0, the routing calculation occurs immediately following the database update.
- Threshold update start. Determines the minimum number of changes to the database before the start of the SPF routing calculations. Default is -1 meaning that the timing of the SPF routing calculation is determined by the configured calculation delay. When set to 0, any database update will cause an SPF routing calculation to occur.
- Threshold update restart. Determines the number of changes to the database before interrupting any running SPF routing calculation and start a new SPF routing calculation. Default is -1 meaning that the no interruptions will occur to SPF routing calculations. When set to 0, a database update will cause any running SPF routing calculation to be restarted.
- Threshold restart limit. Maximum number of restarts allowed before an in-progress SPF routing calculation is completed. Default is 10.

Area authentication

Area authentication supports an additional setting to authenticate sequence number packets (SNPs), including, link state packets, CSNPs, and PSNPs. Default is yes.

Copyright© 2012-2013 Ciena® Corporation

IS-IS supports area and interface authentication to prevent unwanted connections for area and IP interfaces defined by the following:

- Authentication type. Default is none. Optionally, you can set the authentication type to use text or MD5 authentication.
- Password. Default is blank. Required when the authentication type is set to text.
- Message digest. Default is blank. Required when the authentication type is set to MD5.
- Send only. Determines whether authentication occurs only upon sending packets. Default is no.

Interface authentication

Interface authentication supports an additional setting for the router level.

IS-IS supports area and interface authentication to prevent unwanted connections for area and IP interfaces defined by the following:

- Authentication type. Default is none. Optionally, you can set the authentication type to use text or MD5 authentication.
- Password. Default is blank. Required when the authentication type is set to text.
- Message digest. Default is blank. Required when the authentication type is set to MD5.
- Send only. Determines whether authentication occurs only upon sending packets. Default is no.

Step	Action

To control link state packet generation

1 Set the maximum lifetime and refresh interval:

```
isis lsp set isis-instance <isis-instance> [max-lifetime
<NUMBER: 350-65535>] [refresh-interval <NUMBER: 1-65535>]
```

where

isis-instance <isis-instance>	is the isis instance name.
max-lifetime <NUMBER: 350-65535>	is the maximum lifetime of an LSP. Default value is 1200 seconds.
refresh-interval <NUMBER: 1-65535>	is the LSP refresh interval. Default value is 900 seconds.

Copyright© 2012-2013 Ciena® Corporation

2 Set the maximum lifetime and refresh interval to defaults:

```
isis lsp unset isis-instance [max-lifetime <NUMBER: 350-65535>] [refresh-interval <NUMBER: 1-65535>]
```

where

isis-instance <isis-instance>	is the isis instance name.
max-lifetime <NUMBER: 350-65535>	is the maximum lifetime of an LSP. Default value is 1200 seconds.
refresh-interval <NUMBER: 1-65535>	is the LSP refresh interval. Default value is 900 seconds.

To configure Shortest Path First calculations

3 Configure SPF calculation settings:

```
isis spf-calculations set isis-instance <isis-instance> [max-delay <MILLISECONDS>] [threshold-restart-limit <NUMBER: 1-100>] [threshold-update-restart <NUMBER: 1-100>] [threshold-update-start <NUMBER: 1-100>]
```

where

isis-instance <isis-instance>	is the isis instance name.
max-delay <MILLISECONDS>	is the maximum delay (msecs) to start a computation.
threshold-restart-limit <NUMBER: 1-100>	is the maximum number of restarts before an in-progress compution run is completed.
threshold-update-restart <NUMBER: 1-100>	is the minimum number of changes before the restart of a in-progress computation.
threshold-update-start <NUMBER: 1-100>	is the minimum number of changes before the start of computation.

4 Reset SPF calculation settings to default:

```
isis spf-calculations unset isis-instance <isis-instance> [max-delay] [threshold-restart-limit][threshold-update-restart] [threshold-update-start]
```

where

isis-instance <isis-instance>	is the isis instance name.
max-delay	is the maximum delay (msecs) to start a computation.

Copyright© 2012-2013 Ciena® Corporation

where

threshold-restart-limit	is the maximum number of restarts before an in-progress compution run is completed.
threshold-update-restart	is the minimum number of changes before the restart of a in-progress computation.
threshold-update-start	is the minimum number of changes before the start of computation.

To configure an IS-IS area authentication

5 Configure IS-IS area authentication:

```
isis area-authentication set isis-instance set isis-instance <isis-instance> [authentication-type {md5 | text | none}] [password <Password String>][secret <String[62]>] [send-only {yes | no} | snp-authenticate {yes | no}]
```

where

isis-instance <isis-instance>	is the isis instance name.
authentication-type {md5 | text | none}	is the type of authentication.
password <Password String>	is the password.
[secret <String[62]>]	is the encoded password.
send-only {yes | no}	determines whether authentication occurs only upon sending packets. Default is no.
snp-authenticate {yes | no}	is the CSNP,PSNP PDU validation.

Te reset IS-IS area authentication to the deafult values

6 Reset IS-IS area authentication to the default values:

```
isis area-authentication unset isis-instance <isis-instance> [authentication-type] [send-only] [snp-authenticate]
```

where

isis-instance <isis-instance>	is the isis instance name.

Copyright© 2012-2013 Ciena® Corporation

where

authentication-type	is the type of authentication.
send-only	determines whether authentication occurs only upon sending packets. Default is no.
snp-authenticate	is the CSNP,PSNP PDU validation.

To configure IS-IS interface authentication

7 Configure IS-IS interface authentication:

```
isis interface-authentication set ip-interface <ip-interface> [authentication-type {md5 | text | none}] [password <Password String>] [secret <String[62]>] [send-only {yes | no}] level <L1>
```

where

ip-interface <ip-interface>	is the IP interface name.
authentication-type {md5 | text | none}	is the type of authentication.
password <Password String>	is the password.
[secret <String[62]>]	is the encoded password.
send-only {yes | no}	determines whether authentication occurs only upon sending packets. Default is no.
level <L1>	is the routing level.

To reset IS-IS interface authentication to the defaults

8 Reset IS-IS interface authentication to the defaults:

```
isis interface-authentication unset ip-interface <name> [authentication-type] [level] [send-only]
```

—end—

Example

The following example configures an IS-IS area authentication.

```
isis area-authentication set isis-instance test1 authentication-type md5 password myPassword
```

The following example configures an IS-IS interface authentication.

```
isis interface-authentication set l3ip-interface ip-1 authentication-type md5 password myPassword level L1
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 11-14
Displaying IS-IS interface authentication

Display IS-IS interface authentication.

Step	Action
1	Display IS-IS interface authentication: `isis interface-authentication show [ip-interface <name>]` where ip-interface <name> is the IP interface.

—end—

Example

The following example displays IS-IS interface authentication.

```
isis interface-authentication show l3ip-interface ip-1
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 11-15
Associating an IP interface with an IS-IS instance

When you associate an IP interface with the IS-IS instance, it becomes enabled. You can disable or enable individual interfaces and set whether or not to pad hello packets up to the MTU size of the associated IP interface. Default is to use the “padded hello.” Also, you can control the interface Level 1 settings including:

- CSNP interval in seconds. Minimum time between successive CSNPs. Default is 10 seconds.
- Hello multiplier. Determines the hold time for transmitted hello packets. Default is 3.
- Hello timer. Minimum time between successive Hello packets. Default is 10 seconds.
- Link state packet interval. Minimum time between successive link state packets at the level on the circuit. Default is 50 milliseconds.
- Priority. Determines designated router priority. Default is 64. Priority 0 means the interface does not become the designated router.
- Wide metric. Cost of the link. Default is 128.

  Note: In order for higher level protocols, such as RSVP-TE and LDP, to use the route information that IS-IS provides, a loopback interface is required.

Step	Action
1	Associate an IP interface with an IS-IS instance: `isis interface attach ip-interface <interface-name> isis-instance <instance-name> [level <L1 | L2 | L1L2>]`

—end—

Procedure 11-16
Disassociating an IP interface from an IS-IS instance

Disassociate an IP interface from an IS-IS instance.

Step	Action
1	Disassociate an IP interface from an IS-IS instance: `isis interface detach ip-interface <interface-name>`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-17
Enabling or disabling IS-IS for an IP interface

Enable IS-IS for an IP interface.

Disable IS-IS for an IP interface.

Step	Action

To enable IS-IS for an IP interface:

1 Enable IS-IS for an IP interface:

```
isis interface enable ip-interface <interface-name>
```

To disable IS-IS for an IP interface

2 Disable IS-IS for an IP interface:

```
isis interface disable ip-interface <interface-name>
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-18
Configuring IS-IS padded hello for an IP interface

Modify IS-IS padded hello for an IP interface.

Reset IS-IS padded hello to the defautl values for an IP interface.

Step	Action

To modify IS-IS padded hello for an IP interface

1 Modify IS-IS padded hello for an IP interface:

```
isis interface set ip-interface <interface-name> padded-
hello <yes | no>
```

To reset IS-IS padded hello to the default for an IP interface

2 Reset IS-IS padded hello to the default for an IP interface:

```
isis interface unset ip-interface <interface-name>
padded-hello <yes | no>
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-19
Configuring IS-IS interface level attributes

You can

- Modify IS-IS interface level attributes for an IP interface.
- Reset IS-IS interface Level attributes to the default values for an IP interface.

Step	Action

To modify IS-IS interface level attributes for an IP interface

1 Modify IS-IS interface level attributes for an IP interface:

```
isis interface-level set [level <L1 | L2>] [priority
<number>] [hello-timer <secs>] [hello-multiplier
<number>] [wide-metric <number>] [lsp-interval <msecs>]
[csnp-interval <secs>]
```

To reset IS-IS interface level attributes to the default values for an IP interface:

2 Reset IS-IS interface level attributes to the default values for an IP interface:

```
isis interface-level unset ip-interface <interface-name>
[level <L1 | L2>] [priority] [hello-timer] [hello-
multiplier] [wide-metric] [lsp-interval] [csnp-interval
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-20
Enabling or disabling RSVP-TE

RSVP-TE is used for dynamic LSP tunnels. To create an LSP tunnel, the first MPLS device on the path creates an RSVP-TE path message. A label binding for this path is requested and indicates which network layer protocol is to be carried over the path. When the sending device finds a path that either meets the tunnel's QoS requirements, satisfies the policies criteria, or can maximize the use of the network resources, an explicit route is specified. This explicit route can be dynamically changed if the device finds a better route. This event is recorded and the sender device is notified.

The session created, or LSP, carries the load to the destination device along the path. When the destination is reached, a received message is sent back to the sending device, following the path in reverse order, establishing an accurate and effective LSP.

By default, RSVP-TE is disabled globally and for specific IP interfaces. When RSVP-TE is disabled globally, all MPLS ingress tunnels created by RSVP-TE are operationally disabled. RSVP-TE signals over an IP interface when it is enabled.

Step	Action

To enable RSVP-TE globally or for a specific IP interface

1 Enable RSVP-TE globally or for a specific IP interface:

```
rsvp-te enable [ip-interface <interface-name>]
```

where

ip-interface <interface-name>	is the IP interface.

To disable RSVP-TE globally

2 Disable RSVP-TE globally:

```
rsvp-te disable [ip-interface <interface-name>]
```

where

ip-interface <interface-name>	is the IP interface.

—end—

Copyright© 2012-2013 Ciena® Corporation

Example

The following example enables RSVP-TE.

```
> rsvp-te enable
> rsvp-te enable l3ip-interface chassis-1-lpbk
> rsvp-te enable l3ip-interface intf-pe-2
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 11-21
Configuring global RSVP-TE attributes

You can

- set the retry interval
- turn on the retry infinite attribute
- turn off the retry infinite attribute

The retry-infinite attribute must be set to off for the retry interval value to take effect. RSVP-TE attempts to restore the tunnels ten times at the retry interval before it gives up. The retry-interval attribute can only be set if RSVP is globally disabled.

If the retry infinite attribute is turned on, when RSVP tunnels go down because of any reason other than admin = down, RSVP tries to restore the tunnel infinite times. The retry infinite attribute can only be set if RSVP-TE is globally disabled. The default setting is on.

Step	Action

To set the retry interval

1 Set the retry interval:

```
rsvp-te set retry-interval <NUMBER: 3...65>
```

where

retry-interval <NUMBER: 3...65>	is the time interval in seconds that RSVP-TE attempts to restore tunnels. The default setting is 3 seconds.

To turn on the retry infinite attribute

2 Turn on the retry infinite:

```
rsvp-te set retry-infinite on
```

where

retry-infinite on	enables infinite retries for RSVP tunnels

To turn off the retry infinite attribute

3 Turn off the retry infinite:

```
rsvp-te set retry-infinite off
```

where

retry-infinite off	disables infinite retries for RSVP tunnels

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-22
Displaying RSVP-TE configuration and statistics

Display RSVP-TE configuration and statistics.

Step	Action
1	Display RSVP-TE configuration and statistics: `rsvp-te show [statistics]`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-23
Configuring RSVP-TE attributes for IP interfaces

You can configure RSVP-TE attributes for specific IP interfaces, including:

- Advertised-label. Sets the advertised label to support an Implicit Null Label or a non-reserved label (default).
- Hello interval. The RSVP-TE hello message interval ranging from 0-30 seconds. Default is 0, which is disabled.
- Hello tolerance. The RSVP-TE hello tolerance defines number of hello intervals which may pass without receiving a successful Hello message from a partner before the Hello session times out. The range is 0-10, and the default is 3.
- Authentication type. Default is none. Optionally, you can set the authentication type to use MD5 authentication.
- Message digest. An 8-30 character password. Default is blank. Required when the authentication type is set to MD5.

Step	Action
1	Set the RSVP-TE attributes for IP interfaces: `rsvp-te set ip-interface <name> {[advertised-label <implicit-null | non-reserved>] [hello-interval <NUMBER: 0-30>] [hello-tolerance <NUMBER: 0-10>] [authentication-type <md5>] [password <password>]` **—end—**

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-24
Configuring RSVP-TE IP address authentication

In addition to authentication at the IP interface level, RSVP-TE supports authentication at the IP address level and optionally to use MD5 authentication, including the following parameters:

- Authentication type. Default is blank. Optionally, you can set the authentication type to use MD5 authentication.
- Message digest. An 8-30 character password. Default is blank. Required when the authentication type is set to MD5.

In order for IP address level authentication to work, the parent IP interface authentication must also be configured. Also, authentication configuration between two devices affects the status of MPLS tunnels created by RSVP-TE as shown in Table 11-2.

Table 11-2
Authentication Configuration and Tunnel Status

IP address level configured on both	Interface level configured on both	Mesage-digest value on both
Yes	No	Same
Yes	No	Different
Yes	Yes	Same at IP address level. Different at interface level.
Yes	Yes	Different at IP address level. Same at interface level.
Yes	One only	Same at IP address level. Default on interface level.
No	Yes	Same.
No	Yes	Different.
No	One only	Different. Default on IP address and interface levels on non configured device.

You can

- add an IP address
- remove an IP address

Copyright© 2012-2013 Ciena® Corporation

- display IP address entries

Step	Action

To add an IP address

1 Add an IP address:

```
rsvp-te authentication set peer <ip-address>
[authentication-type md5] [password <password>]
```

To remove an IP address

2 Remove an IP address:

```
rsvp-te authentication unset peer <ip-address>
```

To display IP address entries

3 Display IP address entries:

```
rsvp-te authentication show
```

—end—

Example

The following example shows sample output with interface and IP address configuration.

```
> rsvp-te authentication add ip 12.1.0.1 authentication-type md5 message-
digest 123456789
> rsvp-te set l3ip-interface ip-1 authentication-type md5 message-digest
qwertyuiop
> rsvp-te authentication show
|----RSVP NEIGHBOR TABLE-----|
| IpAddress       | AuthType |
|-----------------+----------+
|12.1.0.1         | md5      |
|-----------------+----------+
> rsvp-te interface show
+---------------RSVP INTERFACE TABLE-----------------+--------+----+-----+
| Name                   |Adm|Ope|Type  |Hello|Hello    |Label   |Auth|RSVP |
|                        |   |   |      |     |Tolerance|Type    |    |State|
+------------------------+---+---+------+-----+---------+--------+----+-----+
|ip-1                    |En |En |Bcast |Di   |3        |Non-Rsvd|md5 |En   |
|LBK                     |En |En |Lbk   |Di   |3        |Non-Rsvd|none|En   |
+------------------------+---+---+------+-----+---------+--------+----+-----|
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 11-25
Configuring RSVP-TE paths

The system supports a maximum of 256 paths with 16 hops per path.

Step	Action

To create an RSVP-TE path

1 Create an RSVP-TE path:

```
rsvp-te path create rsvp-path <name>
```

To delete an RSVP-TE path

2 Delete an RSVP-TE path:

```
rsvp-te path delete rsvp-path <name>
```

To set the hops for an RSVP-TE path

3 Set the hops for an RSVP-TE path:

```
rsvp-te path set rsvp-path <name> index <number(1-100)>
ip <ip-address> [hop-type <strict | loose>]
```

To remove a hop from the path

4 Remove a hop from the path:

```
rsvp-te path unset rsvp-path <name> index
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-26
Displaying all or specific paths

Display all or specific paths.

Step	Action
1	Display all or specific paths: `rsvp-te path show [rsvp-path <name>]`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-27
Running traceroute and ping for RSVP-TE tunnels

You can generate traceroute and ping packets for testing RSVP-TE tunnels.

Step	Action

To run a traceroute

1 Run a traceroute:

```
mpls encap-tunnel traceroute rsvp-lsp
<MplsDynamicEncapTunnel> [timeout <NUMBER: 1000-10000>]
[ttl <NUMBER: 1-30>]
```

To run a ping

2 Run a ping:

```
mpls encap-tunnel ping rsvp-lsp
<TunnelDynamicEncapName[31]> [count <NUMBER: 1-100>]
[packet-size <NUMBER: 96-1464>] [timeout <NUMBER: 1000-
10000>] [ttl <NUMBER: 1-255>]
```

—end—

Example

The following example shows sample output for a traceroute operation.

```
> mpls encap-tunnel encap traceroute rsvp-lsp DynamicLSPA


  0 10.10.50.2 MRU 1500 Label: 28871


! 1 10.10.20.20 5 ms
```

The following example shows sample output for a ping operation.

```
> mpls encap-tunnel ping rsvp-lsp DynamicLSPA count 100
packet-size 1464


!!!!!!!!!
--------------- Statistics ---------------


100 packets transmitted, 9 packets received


round-trip (ms) min/avg/max = 10/20/89
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 11-28
Configuring tunnels

MPLS tunnels form the LSPs to transport MPLS traffic. The system supports three types of tunnels depending upon the role of the switch in the MPLS cloud:

- For a switch operating as an ingress LER, ingress tunnels are created to encapsulate incoming traffic with the tunnel label. Ingress tunnels can be created dynamically, where RSVP-TE is used for signalling and label exchange, or statically, where labels are explicitly defined.
- For a switch operating as an egress LER, egress tunnels decapsulate traffic to remove the tunnel label and transmit it out of the MPLS cloud.

  Note: In configurations using RSVP-TE, the egress tunnels are dynamically created.

- For a switch operating as an LSR, transit tunnels form the parts of the LSP path to carry the encapsulated traffic between the LERs.

  Note: In configurations using RSVP-TE, the transit tunnels are dynamically created.

The system supports the following:

- 500 ingress tunnels. Ciena recommends a maximum of 400 ingress tunnels per port.
- 500 egress tunnels. Ciena recommends a maximum of 400 egress tunnels per port.
- 1000 transit tunnels. Ciena recommends a maximum of 400 transit tunnels per transit pair port.
- 1000 tunnels per chassis, where the number of tunnels is the sum of egress tunnels, ingress tunnels, and transit tunnels
- 1000 VCs (PWs)
- 250 VS (VPLS)

  Note: The system does not block users from configuring more tunnels than the maximum recommended. If the warning is ignored and users proceed to configure beyond the recommended numbers, the system might become unstable.

Copyright© 2012-2013 Ciena® Corporation

Creating and modifying ingress tunnels

Dynamic, static ingress, and static transit tunnels share common attributes as listed in Table 11-3.

Table 11-3
Common tunnel attributes

Attribute	Description
Time To Live (TTL) Policy	When set to inherit, the value of the TTL in the tunnel label is inherited from the IP TTL of customer frame. When set to fixed, the value of the TTL in the tunnel labelis set according to the fixed TTL (fixed-ttl) value. The default (and only supported) policy is fixed.
Fixed TTL	Applicable when the TTL policy is set to fixed. The default is 255.

Dynamic and static ingress tunnels share additional common attributes as listed in Table 11-4.

Table 11-4
Commion ingress tunnel attributes

Attribute	Description
Destination IP address	Sets the destination IP address for the tunnel.
Backup tunnel	Ingress tunnel to protect the primary tunnel.
Reversion hold time	Sets the amount of time in seconds to wait before switching from the backup to the primary tunnel after the fault on the primary is cleared (tunnel reversion). Default is 300 seconds. Applicable when tunnel reversion is turned on.
Tunnel reversion	Indicates whether or not to switch from the backup to the primary tunnel after the fault on the primary is cleared.

Copyright© 2012-2013 Ciena® Corporation

Dynamic ingress tunnels have additional attributes as listed in Table 11-5.

Table 11-5
Dynamic ingress additional tunnel attributes

Attribute	Description
Setup priority	Setup priority is compared with the hold priority of existing links to determine if preemption of those links should be allowed. The setup of new tunnels requires this contention system when resources for new services have been exceeded. The default setup priority is 7, which is the lowest priority.
Hold priority	Setup priority is compared with the hold priority of existing links to determine if preemption of those links should be allowed. The setup of new tunnels requires this contention system when resources for new services have been exceeded. Default hold priority is 0.
Record route	For fastest reroute protection, the record route needs to be on. The default is on. Fast ReRoute protection type attribute: protection-type <link | node>
Path	Explicit RSVP-TE path.
Bandwidth profile	Associates the tunnel with a bandwidth profile that defines bandwidth and burst size signaling information to be carried in RSVP-TE path messages for interoperability with other vendor core routers. You can create up to 64 bandwidth profiles with the following attributes: • Bandwidth. Required attribute specifying the rate in kilobits per second (kbps). • Burst. Optional attribute specifying the burst size in kilobytes. Default is 128 kilobytes.

Static ingress tunnels have additional attributes as listed in Table 11-6.

Table 11-6
Static ingress tunnel additional attributes

Attribute	Description
Next hop IP address	IP address of the next hop device.
Label	Tunnel label selected from the configured static label range.

- Fast-reroute protection. Fast reroute provides a way to dynamically establish backup tunnels, or detour paths, in order to reduce packet loss on the LSP. If there is a link or node failure, fast reroute RSVP-TE tunnels automatically establish local detours around the failed link or node and then redirect user traffic to the detour. Options are link-protect or node-protect. The default is link-protect. These tunnels are counted in the

Copyright© 2012-2013 Ciena® Corporation

maximum allowed dynamic encapsulation. So, if all dynamic encapsulation tunnels are configured with fast-reroute protection, the number should be limited to 128 to allow for 128 fast-reroute tunnels.

You can

- create a tunnel bandwidth profiie
- delete a tunnel bandwidth profile

Step	Action

To create a tunnel bandwidth profile

1 Create a tunnel bandwidth profile:

```
mpls tunnel-bandwidth-profile create bandwidth-profile
<profile-name> bandwidth <NUMBER: 1000-10000000> [burst
<NUMBER: 12-256>]
```

To delete a tunnel bandwidth profile

2 Delete a tunnel bandwidth profile:

```
mpls tunnel-bandwidth-profile delete bandwidth-profile
<profile-name>
```

—end—

Procedure 11-29
Displaying tunnel bandwidth profiles

Display tunnel bandwidth profiles.

Step	Action
1	Display tunnel bandwidth profiles: `mpls tunnel-bandwidth-profile show`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-30
Configuring dynamic ingress tunnels

You can

- create a dynamic ingress tunnel
- modify a dynamic ingress tunnel

Step	Action

To create a dynamic ingress tunnel

1 Create a dynamic ingress tunnel:

```
mpls tunnel create rsvp-ingress <lsp-name> dest-ip <ip-
address> [frame-cos-policy <fixed | rcos-mapped>] [frame-
cos-map <fcos-map-name>] [fixed-tc <0-7>] [setup-priority
<0-7>] [hold-priority <0-7>] [record-route <on | off>]
[protection-type <link | node>] [path <path-name>] [ttl-
policy <fixed | inherit>] [fixed-ttl <1-255>] [bandwidth-
profile <profile-name>] [backup-tunnel <rsvp-ingress-
tunl-list>] [tunnel-reversion <on | off>] [reversion-
hold-time <0-3600>]To modify a dynamic ingress tunnel
```

2 Modify a dynamic ingress tunnel:

```
mpls tunnel set rsvp-ingress <lsp-name> [frame-cos-policy
< fixed | rcos-mapped >] [frame-cos-map <fcos-map-name>]
[fixed-tc <0-7>] [setup-priority <0-7>] [hold-priority
<0-7>] [record-route <on | off>] [path <path-name>] [ttl-
policy <fixed | inherit>] [fixed-ttl <1-255>] [reversion-
hold-time <0-3600>] [tunnel-reversion <on | off>]
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-31
Configuring static ingress tunnels

You can

- create a static ingress tunnel
- modify a static ingress tunnel
- reset static ingress tunnel attributes to default values

Step	Action

To create a static ingress tunnel

1 Create a static ingress tunnel:

```
mpls tunnel create static-ingress <lsp-name> dest-ip <ip-address> next-hop-ip <ip-address> label <NUMBER: 0-1048575> [frame-cos-policy < fixed | rcos-mapped >] [frame-cos-map <fcos-map-name>] [fixed-tc <0-7>] [ttl-policy <fixed | inherit>] [fixed-ttl <1-255>] [backup-tunnel <static-ingress-tunl-list>] [reversion-hold-time <0-3600>] [tunnel-reversion <on | off>]To modify a static ingress tunnel
```

2 Modify a static ingress tunnel:

```
mpls tunnel set static-ingress <lsp-name> [label <0-1048575>] [frame-cos-policy < fixed | rcos-mapped >] [frame-cos-map <fcos-map-name>] [fixed-tc <0-7>] [ttl-policy < fixed | inherit >] [fixed-ttl <1-255>] [reversion-hold-time <0-3600>] [tunnel-reversion <on | off>]
```

To reset static ingress tunnel attributes to default values

3 Reset static ingress tunnel attributes to default values:

```
mpls tunnel unset static-ingress <lsp-name> [frame-cos-policy] [frame-cos-map] [fixed-tc] [ttl-policy] [fixed-ttl] [reversion-hold-time] [tunnel-reversion]
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-32
Configuring static egress tunnels

In a static deployment, an egress tunnel must be created to decapsulate the tunnel headers from the egress traffic. A static egress tunnel is defined by the source IP address and the unique MPLS label from the range of configured static labels.

You can

- create a static egress tunnel
- delete a static egress tunnel
- modify the label for a static egress tunnel

Step	Action

To create a static egress tunnel

1 Create a static egress tunnel:

```
mpls tunnel create static-egress <lsp-name> label <0-
1048575> [src-ip <ip-address>]To delete a static egress
tunnel
```

2 Delete a static egress tunnel:

```
mpls tunnel delete static-egress <StaticEgressTunnelList>
```

To modify the label for a static egress tunnel

3 Modify the label for a static egress tunnel:

```
mpls tunnel set static-egress <lsp-name> [label <0-
1048575>
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-33
Configuring static transit tunnels

Static transit tunnels support configurable in and out label settings in addition to the common TTL policy and Fixed TTL attributes.

You can

- create a static transit tunnel
- delete a static transit tunnel
- modify the label for a static transit tunnel

Step	Action

To create a static transit tunnel

1 Create a static transit tunnel:

```
mpls tunnel create static-transit <name> next-hop-ip <ip>
{in-label <16-1048575>} {out-label <16-1048575>} [ttl-
policy <fixed | inherit> default: inherit] [fixed-ttl <1-
255> default: 255]
```

To delete a static transit tunnel

2 Delete a static transit tunnel:

```
mpls tunnel delete static-transit <name>
```

To modify the label for a static transit tunnel

3 Modify the label for a static transit tunnel:

```
mpls tunnel set static-transit <name> {[in-label <16-
1048575>], [out-label <16-1048575>], [ttl-policy <fixed |
inherit> default: inherit], [fixed-ttl <1-255> default:
255]}
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-34
Enabling and disabling tunnels

You can administratively enable a tunnel to potentially change the operational state to up. You can administratively disable a tunnel to change the operational state to down.

You can

- enable a tunnel
- disable a tunnel

Step	Action

To enable a static or dynamic ingress tunnel

1 Enable a static or dynamic ingress tunnel:

```
mpls tunnel enable lsp <lsp-name>
```

where

lsp <lsp-name> is the tunnel to be enabled.

To disable a static or dynamic ingress tunnel

2 To disable a static or dynamic ingress tunnel:

```
mpls tunnel disable lsp <lsp-name>
```

where

lsp <lsp-name> is the tunnel to be disabled.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-35
Deleting ingress tunnels

You can delete a manually created static or dynamic ingress tunnel.

Step	Action
1	Delete a static or dynamic ingress tunnel: `mpls tunnel delete lsp <lsp-name>` where lsp <lsp-name> is the tunnel to be deleted.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-36
Displaying tunnels

You can display a list of tunnels, including the type, index, name, destination IP address, label, administrative state, and operational state. Optionally, you can display details about a specific tunnel. Also, you can filter to display a list of tunnels by tunnel configuration (static or dynamic), type (ingress or egress), state (up or down), or a specific static egress tunnel.

Step	Action

To display all or a specific tunnel

1 Display all or a specific tunnel:

```
mpls tunnel show lsp <lsp-name>
```

where

lsp <lsp-name> is the tunnel to be displayed.

To display a filtered list of tunnels

2 Display a filtered list of tunnels:

```
mpls tunnel show {[static-lsp | rsvp-lsp] [type < ingress
| egress>] [state <up | down>] [static-egress-tunnel
<static-tunnel-name>]}
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-37
Configuring LDP globally

LDP is required for dynamic MPLS deployments for signaling virtual circuits. Sessions are established using UDP port 646 to a peer IP address that can be directly connected or several hops away.

By default, LDP is globally disabled. For dynamic deployments, it must be globally enabled. Configurable LDP attributes include the hello hold time and keep alive interval used when establishing the LDP targeted session.

You can

- enable LDP globally
- disable LDP globally
- modify global LDP attributes

Step	Action

To enable LDP globally

1 Enable LDP globally:

```
ldp enable
```

To disable LDP globally

2 Disable LDP globally:

```
ldp disable
```

To modify global LDP attributes

3 Modify global LDP attributes:

```
ldp set {[hello-hold-time <NUMBER: 1-65535>], [keep-
alive-hold-time <NUMBER: 1-65535>]}
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-38
Displaying LDP configuration

Display LDP configuration.

Step	Action
1	Display LDP configuration: `ldp show [adjacency] [sessions] [timers]`

—end—

Example

The following example shows a global configuration.

```
> ldp show
+------------------- LDP GLOBAL CONFIG ---------------+
|         Parameter          |           Value         |
+----------------------------+-------------------------+
| LDP Admin State            | Disabled                |
| LDP Oper  State            | Disabled                |
+----------------------------+-------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 11-39
Configuring LDP authentication

LDP supports authentication to prevent unwanted connections. To configure authentication, add IP entries and passwords for Peer Nodes that are allowed to establish sessions. By default, LDP supports MD5 based authentication.

You can

- add an IP entry and password
- remove an IP entry and password
- display IP entries with an encoded password

Step	Action

To add an IP entry and password

1 Add an IP entry and password:

```
ldp authentication set {peer <IPAddress>}
{authentication-type <md5>} {password
<PasswordString[31]>} {secret <SecretString[62]>}
```

To remove an IP entry and password

2 Remove an IP entry and password:

```
ldp authentication unset {peer <IPAddress>}
```

To display IP entries with an encoded password

3 Display IP entries with an encoded password:

```
ldp authentication show
```

—end—

Example

The following example adds an IP entry and password.

```
ldp authentication set peer 1.2.3.4 password myPassword
ldp authentication show
+------------ LDP Authentication Configuration Summary -------------------+
|   Router Id    |   Encoded Password                                      |
+----------------+---------------------------------------------------------+
|1.2.3.4         |ffe5109e033a3716e94f                                     |
+----------------+---------------------------------------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 11-40
Configuring virtual circuits

When configuring primary virtual circuits, follow these configuration rules:

- All virtual circuit names must be unique.
- A virtual circuit cannot be deleted unless it is detached from the virtual switch.
- Dynamic virtual circuits: For a given peer node, PW-ID of virtual circuits must be unique.
- Static VCs: Ingress label is unique not only across virtual circuits but also tunnels.

For dual-homed protection of an MTU-s, the secondary virtual circuit is defined by setting the peer IP address, tunnel, primary PW-ID, and secondary PW-ID. All other attributes are inherited from the primary.

By default, the primary virtual circuit is the active virtual circuit. If the primary fails, the secondary virtual circuit becomes the active virtual circuit. The active virtual circuit (whether it is the primary or secondary) remains active unless it fails or it can be manually switched over.

Step	Action

To create a dynamic virtual circuit

1 Create a dynamic virtual circuit:

```
mpls l2-vpn create dynamic-vc <name> pw-id <1-4294967295>
peer <IP Address> tunnel <lsp-name> [pw-type <eth-
raw|eth-tagged>] [mtu<1500-9128>] [status-tlv <on|off>]
[pw-mode <mesh| spoke>][service-delimiter-vid <NUMBER: 1-
4094>] [service-delimiter-tpid <8100|9100|88A8>]
```

To modify mode, customer name, service delimiter VID, or service delimiter TPID for a dynamic virtual circuit:

2 Modify mode, customer name, service delimiter VID, or service delimiter TPID for a dynamic virtual circuit:

```
mpls l2-vpn set dynamic-vc <name> [tunnel <lsp-name>]
[pw-mode <mesh| spoke>] [customer <name>] [service-
delimiter-vid <NUMBER: 1-4094>] [service-delimiter-tpid
<8100|9100|88A8>]
```


Copyright© 2012-2013 Ciena® Corporation

To set mode, customer name, service delimiter VID, or service delimiter TPID to the default values

3 Set mode, customer name, service delimiter VID, or service delimiter TPID to the default values:

```
mpls l2-vpn unset dynamic-vc <name> [pw-mode]
[customer][service-delimiter-vid] [service-delimiter-
tpid]
```

To create a static virtual circuit

4 Create a static virtual circuit:

```
mpls l2-vpn create static-vc <name> pw-id <pw-id> peer
<peer-ip> ingress-label <> egress-label <> [tunnel <lsp-
name>] [pw-type <raw | tagged>] [status-tlv <on | off>]
[mtu <1500-9128>] [pw-mode <mesh|spoke>] [service-
delimiter-vid <NUMBER: 1-4094>] [service-delimiter-tpid
<8100|9100|88A8>]
```

To modify mode, customer name, service delimiter VID, or service delimiter TPID for a static virtual circuit

5 Create a dynamic virtual circuit:

```
mpls l2-vpn set static-vc <name> [tunnel <lsp-name>] [pw-
mode <mesh| spoke>] [customer <name>] [service-delimiter-
vid <NUMBER: 1-4094>] [service-delimiter-tpid
<8100|9100|88A8>]
```

To set mode, customer name, service delimiter VID, or service delimiter TPID for a static virtual circuit to the default values

6 Set mode, customer name, service delimiter VID, or service delimiter TPID for a static virtual circuit to the default values:

```
mpls l2-vpn unset static-vc <name> [pw-mode] [customer]
[service-delimiter-vid] [service-delimiter-tpid]
```

To create a dynamic protection virtual circuit

7 Create a dynamic protection virtual circuit:

```
mpls l2-vpn protection create dynamic-vc <name>
secondary-pw-id <secondary-pw-id> primary-pw-id
<primary-pw-id> peer <peer-ip> tunnel <lsp-name>
```

To manually switch over to the protection circuit

8 Manually switch over to the protection circuit:

```
mpls l2-vpn protection switchover dynamic-vc <name>
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-41
Displaying virtual circuits

You can display all static and dynamic virtual circuit types with the name, logical ID, tunnel group ID, administrative state, operational state, whether it is active or not, ingress label, egress label, and virtual switch ID. Optionally, you can display detailed configuration for a specific virtual circuit.

Step	Action
1	Display a virtual circuit: `mpls l2-vpn show [vc <vc-name>]` where vc <vc-name> is the virtual circuit to display.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-42
Disabling and enabling virtual circuits

When a virtual circuit is disabled, the operational state is down. When it is enabled, the operational state is up.

Step	Action

To disable a virtual circuit

1 Disable a virtual circuit:

```
mpls l2-vpn disable [vc <vc-name>]
```

where

vc <vc-name> is the virtual circuit to disable.

To enable a virtual circuit

2 Enable a virtual circuit:

```
mpls l2-vpn enable [vc <vc-name>]
```

where

vc <vc-name> is the virtual circuit to enable.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-43
Deleting a virtual circuit

Delete a virtual circuit when it is no longer in use.

Step	Action
1	Delete a virtual circuit: `mpls l2-vpn delete [vc <vc-name>]` where vc <vc-name> is the virtual circuit to delete.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-44
Configuring label ranges

You can configure

- static label ranges
- dynamic label ranges

MPLS label ranges can be modified at any time, however, the chassis must be rebooted in order for the new range to become operational. Also, label ranges between static and dynamic cannot overlap.



CAUTION
Tunnels and VCs Could be Removed from Configuration
If there are any tunnels or virtual circuits configured to use labels outside of the new range, they will be removed from the configuration upon reboot.

Egress labels for static tunnels and virtual circuits can be any valid label between 16-1048575. Ingress labels can only be from the specified MPLS static label range.

The dynamic label range should be same on both ends of a tunnel or virtual

Step	Action

To configure static label ranges

1 Configure static label ranges:

```
mpls static-label-range set min-label <16-1048575> max-label <16-1048575>
```

To configure dynamic label ranges

2 Configure dynamic label ranges:

```
mpls dynamic-label-range set min-label <16-1048575> max-label <16-1048575>
```

—end—

Example

The following example sets the static label range with a minimum label value of 30 and a maximumlabel value of 1023.

```
mpls static-label-range set min-label 30 max-label 1023
```

The following example sets the dynamic label range with a minimum label value of 2048 and a maximum label value of 131071.

Copyright© 2012-2013 Ciena® Corporation

```
mpls dynamic-label-range set min-label 2048 max-label
131071
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 11-45
Displaying label ranges

You can display

- static label ranges
- dynamic label ranges

Step	Action

To display static label ranges

1 Display static label ranges:

```
mpls static-label-range show
```

To display dynamic label ranges

2 Display dynamic label ranges:

```
mpls dynamic-label-range show
```

—end—

Example

The following example shows sample output for the mpls static-lable-range show command.

```
> mpls static-label-range show
+--------------- MPLS Static Label Range -------------+
| MPLS Min Static Label      | 30                     |
| MPLS Max Static Label      | 1023                   |
+----------------------------+------------------------+
```

The following example shows sample output for the mpls dynamic-label-range show command.

```
> mpls dynamic-label-range show
+-------------- MPLS Dynamic Label Range -------------+
| MPLS Min Dynamic Label     | 2048                   |
| MPLS Max Dynamic Label     | 131071                 |
+----------------------------+------------------------+
```

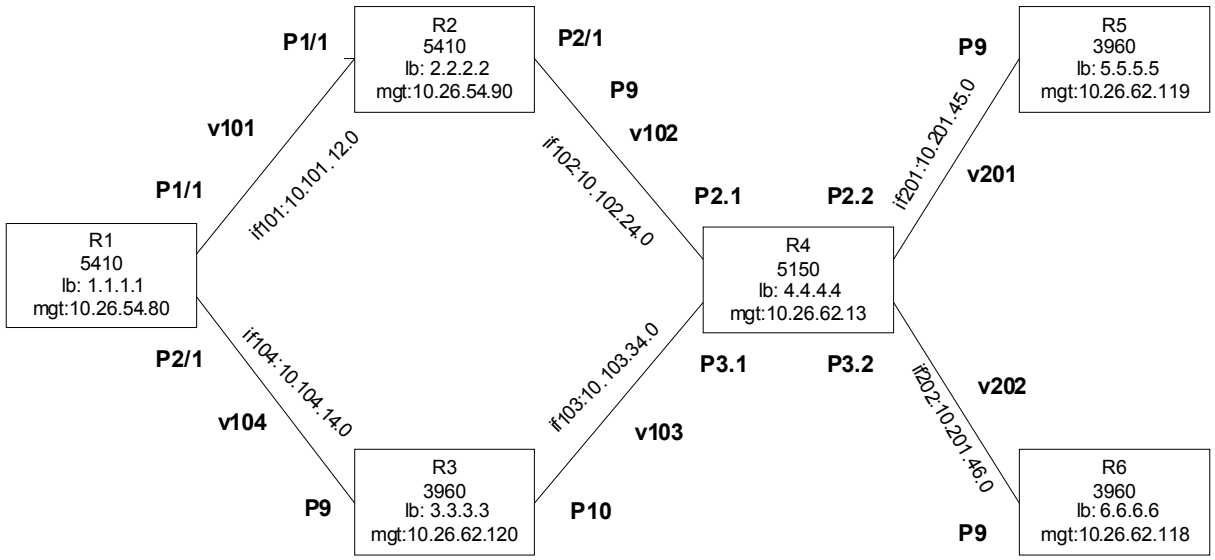

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-46
Base MPLS configuration (LSR) example

This section reviews the common set of options required for any L2VPN service configuration. The base configuration allows the switch to function as an LSR and builds the foundation for LER functionality. This section will setup the L3 interfaces, IGP, and Signaling protocols used by MPLS.

Figure 11-8 shows a sample physical topology.

Figure 11-8
Sample Physical Topology



To configure an LSR:

Step	Action
1	Create a VLAN for each L3 interface / physical port pair. `vlan create vlan <vlan_#>`
2	Add the L3 Interface VLAN to the associated physical port. `vlan add vlan <vlan_#> port <port_#>`
3	Optional: (Only required in a ring topology) Remove VLAN 1,127 from each of the physical ports. `vlan remove vlan 1,127 port <port_#>`
4	Disable RSTP. `rstp disable`

Copyright© 2012-2013 Ciena® Corporation

5 Create the L3 and Loopback interfaces.

```
interface create loopback <name> ip <ip>
```

```
interface create ip-interface <name> ip <ip> subnet
<subnet> vlan <vlan> ip-forwarding <on|off>
```

6 Create IGP instance.

```
isis instance create isis-instance <instance_name> level
L1 area <area_#>
```

7 Attach each L3 interface to the IGP instance.

```
isis interface attach ip-interface <name> isis-instance
<name> level L1
```

8 Enable RSVP globally and on each of the interfaces previously created.

```
rsvp-te enable
```

```
rsvp-te enable ip-interface <name>
```

9 Enable LDP globally.

```
ldp enable
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Example

The following sample configuration file segment shows the base LSR configuration for the 3960 named R3 in the network topology diagram shown in Figure 11-8.

```
                !----------------------------------------------------
                !     General Device Config
                !----------------------------------------------------
system set host-name R3
                ! VLANs 1,127 have been removed from all ports in the base configuration.
                This eliminates the possibility of L2 loops which allows RSTP to be disabled
                globally.
rstp disable
                !
interface create loopback ip 3.3.3.3
                !
isis instance create isis-instance region1 level L1 area 46.0001
                !----------------------------------------------------
                !     Interface if104
                !----------------------------------------------------
                !A different VLAN is used for each L3 Interface / physical port combination to
                prevent a flood domain from forming.
vlan create vlan 104
vlan add vlan 104 port 9
vlan remove vlan 1,127 port 9
                !The L3 interface named if104 is associated with VLAN 104 to coincide with
                port 9 above. IP Forwarding is also enabled to allow routing between L3
                interfaces.
interface create ip-interface if104 ip 10.104.13.3 subnet 255.255.255.0 vlan
104 ip-forwarding on
                !Each L3 interface is attached to the IGP instance to signal that it should be
                included in routing updates and to enable IGP signaling for that link.
isis interface attach ip-interface if104 isis-instance region1 level L1
                !RSVP provides the signaling and label exchange mechanism that will be used
                to generate the LSP. It must be enabled both globally and on each interface.
rsvp-te enable ip-interface loopback
rsvp-te enable
rsvp-te enable ip-interface if104
                !LDP provides the signaling and label exchange mechanism that will be used
                to generate the Pseudowire. LDP is enabled globally.
                ldp enable
                !----------------------------------------------------
                !     Interface if103
                !----------------------------------------------------
vlan create vlan 103
vlan add vlan 103 port 10
vlan remove vlan 1,127 port 10
                !
interface create ip-interface if103 ip 10.103.34.3 subnet 255.255.255.0 vlan
103 ip-forwarding on
                !
isis interface attach ip-interface if103 isis-instance region1 level L1
                !
rsvp-te enable ip-interface if103
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 11-47
Allocating resources for an MPLS management VS (3916, 3930 and 3931 platforms)

On the 5150 platform, sufficient resources are available for all feature sets and no special treatment is required if the default configuration of the resource manager is used.

On the 3916, 3930 and 3931 platforms, support for an MPLS management VS requires adjustment of the default resource allocation:

- If an attempt is made to use the “interface remote set vs <vs-name>” command before these resources have been allocated on the 3916/30/31 platforms, an error message will be displayed and the command will be refused.
- If an attempt is made to remove or reduce the resources allocated to the transport-oam feature set while the remote interface is associated with a VS, the command will be refused.

Step	Action

1	Allocate resources to the transport-oam feature set:
	`resource-manager pool set feature transport-oam resource classifier count 256`
	`resource-manager pool set feature transport-oam resource counter count 256`

2	Deallocate resources from another feature (do one of the following two options):
	• Reduce resources from the traffic-profiling feature:
	`resource-manager pool set feature traffic-profiling resource classifier count 768`
	`resource-manager pool set feature traffic-profiling resource counter count 768`
	• Eliminate resources from either the broadcast-containment or vc-statistics feature:
	`resource-manager pool set feature <feature-name> resource classifier count 0`
	`resource-manager pool set feature <feature-name> resource counter count 0`
	where
	feature <feature-name> — either “broadcast-containment” or “vc-statistics”

Copyright© 2012-2013 Ciena® Corporation

3 Save the configuration and reboot:

```
config save
reboot now
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-48
Creating an MPLS management VS

Create an MPLS management VS by creating an MPLS VS and associating it with the remote interface.

Note: In order to carry remote interface traffic over an MPLS tunnel, the remote-interface is associated with a Virtual Switch (VS). Management access to the switch can then be gained from any of the members of this VS, including Attachment Circuit (AC) members. Thus, if ACs exist on the VS that is associated with the Remote Interface, customers could obtain management access to the node.

In order to prevent this, create an MPLS VS specifically for use for in-band management, and DO NOT attach customer ACs to that VS.

Note: For 3916, 3930 and 3931 platforms ensure that resource allocations have been adjusted. See "Allocating resources for an MPLS management VS (3916, 3930 and 3931 platforms)" on page 11-81.

Step	Action

1 Create the virtual switch (VS).

```
virtual-switch ethernet create vs <vs> mode vpls
```

where

vs <vs> is the name of the MPLS VS

2 Attach the VS to the MPLS virtual circuit (VC):

```
virtual-switch ethernet attach vs <vs> mpls_vc <vc>
```

where

vs <vs> is the name of the MPLS VS

vc <vc> is the name of the MPLS VC

3 Optionally, add the VS to port and VLAN:

```
virtual-switch ethernet add vs <vs> {port <port>} {vlan
<VLAN list>} {ip-interface <interface-name>} {encap-cos-
policy <dot1dpri-inherit | fixed | ip-prec-inherit |
phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-
```


Copyright© 2012-2013 Ciena® Corporation

```
dot1dpri <NUMBER: 0-7>} [statistics <on | off>]
[translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100
| 9100 | 88A8>}
```

where

vs <vs> is the name of the MPLS VS

4 Associate the VS with the remote interface:

```
interface remote set vs <ve>
```

where

vs <vs> is the name of the MPLS VS

Note: Associating an MPLS VS with the remote interface does not change the basic properties of that VS. Support for features such as Traffic Profiling, L2-CFT, COS Mapping is unchanged.

—end—

If the CLI session in use is connected to the remote interface, the user will, of course, lose access via that current session as the remote interface is reconfigured – and will have to initiate a new session over the newly configured VS.

At the node where VLAN based management transitions to inband MPLS based management, the VLAN based management traffic must be added to the MPLS VS as an EVPL AC (where the VLAN used is the management VLAN)

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-49
Displaying remote interface configuration

Use the “interface remote show” command to identify the Management Domain used for connectivity (either a VLAN or a VS).

Step	Action
1	Display information about the remote interface: `interface remote show` In the example output below, the remote interface connectivity is provided by an MPLS management VS named “`MyMngmt1`”

```
+---------------------------------- INTERFACE STATE -------------------------------+
| Parameter               | Value                              | Source    | State      |
+-------------------------+------------------------------------+-----------+------------+
| Name                    | remote                             |           |            |
| Index                   | 15                                 |           |            |
| Admin State             | Enabled                            |           |            |
| Oper State              | Enabled                            |           |            |
| MAC Address             | 00:02:5a:01:c5:4f                  |           |            |
| Management Domain       | VS MyMngmt1                        |           |            |
| Priority                | 7                                  |           |            |
| MTU                     | 1500                               |           |            |
+-------------------------+------------------------------------+-----------+------------+
| IPv4 Oper addr/mask     | 192.168.50.1/24                    | Manual    | PREFERRED  |
| IPv4 Broadcast Address  | 192.168.50.255                     | Manual    | PREFERRED  |
+-------------------------+------------------------------------+-----------+------------+
| IPv6 Oper addr/mask     | fe80::202:5aff:fe01:c54f/64        | Internal  | PREFERRED  |
+-------------------------+------------------------------------+-----------+------------+
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-50
Deleting an MPLS management VS

A VS that is currently associated with the remote interface cannot be deleted until it is removed from the remote interface by using the “interface remote set” command which attaches an alternative management VS or management VLAN.

Step	Action
1	Attach an alternative management VS or management VLAN. For example: `interface remote set vlan 127`
2	Detach all member MPLS VCs: `virtual-switch ethernet detach vs <vs_name> mpls_vc <vc_name>` where vs <vs-name> is the name of the MPLS VS vc <vc-name> is the name of the MPLS VC
3	Delete the management virtual switch (VS). `virtual-switch ethernet delete vs <vs_name>` where vs <vs-name> is the name of the MPLS VS

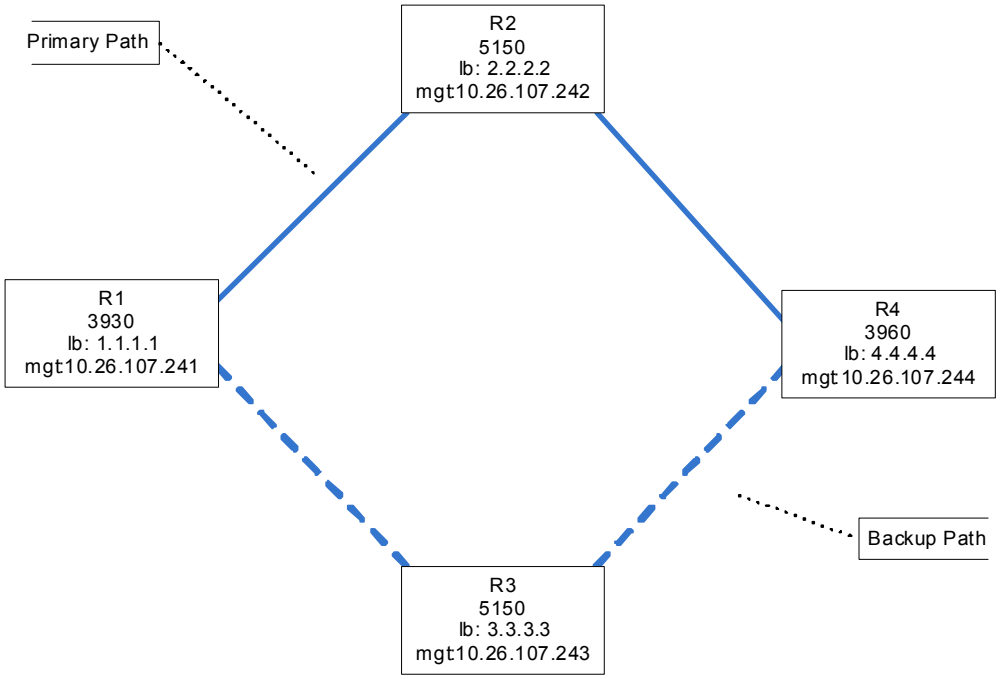
—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 11-51
VPWS configuration example

VPWS provides point-to-point connectivity between two remote Local Area Networks (LANs). In this example protection switching is performed through tunnel redundancy.

Figure 11-9
VPWS Topology



To configure the VPWS service:

Step	Action
1	Create a virtual switch. `virtual-switch ethernet create vs <vs> mode <vpws|vpls>` where vs <vs> is the virtual switch name. mode <vpws|vpls> is the capability of the virtual-switch.

Copyright© 2012-2013 Ciena® Corporation

2 Attach the UNI port(s) to the virtual switch.

```
virtual-switch ethernet add vs <vs> {port <port>} {vlan
<VLAN list>} {ip-interface <interface-name>} {encap-cos-
policy <dot1dpri-inherit | fixed | ip-prec-inherit |
phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-
dot1dpri <NUMBER: 0-7>} [statistics <on | off>]
[translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100
| 9100 | 88A8>}
```

where

vs <vs>	is the virtual switch name.
port <Port list>	is the subscriber port list.
vlan <VLAN list>	is the subscriber VLAN ID.
service-vlan-tpid <8100 | 9100 | 88A8>	is the service VLAN TPID value. The service VLAN TPID is valid only for MPLS-enabled virtual-switches. It denotes the TPID used for frames that are service tagged. When an egress port is in NNI mode (no virtual switch memberships for that port), Ethernet frames egressing the port can have their TPID stamped on a per-VLAN basis with one of the following TPID values: • 0x8100 (802.1Q): to identify the egress frame as single tag VLAN frame • 0x9100 (Q-in-Q): to identify the egress frame as a double-tagged frame • 0x88a8 (Provider Bridge): to identify the egress frame as a PBB-TE frame

3 Create a primary and backup path.

```
rsvp-te path create rsvp-path <rsvp-path>
```

where

rsvp-path <rsvp-path>	is the RSVP path to create.

```
rsvp-te path set rsvp-path <rsvp-path> index
<NUMBER: 1-16> ip <IP address> hop-type <strict|loose>
```

where

rsvp-path <rsvp-path>	is the RSVP path to create.
index <NUMBER: 1-16>	is the position of the node in the rsvp path list.
ip <IP address>	is the IP address.
hop-type <strict|loose>	is the hop type.

Copyright© 2012-2013 Ciena® Corporation

4 Create a primary and backup tunnel and apply the previously created path.

```
mpls tunnel create rsvp-ingress <rsvp-ingress> dest-ip
<IP address> path <MPLS Rsvp Path> [backup-tunnel <MPLS
ingress primary tunnel>]
```

where

rsvp-ingress <rsvp-ingress>	is the RSVP tunnel name.
dest-ip <IP address>	is the tunnel destination IP address.
path <MPLS Rsvp Path>	is the RSVP path name.
[backup-tunnel <MPLS ingress primary tunnel>]	is the list of primary tunnels.

5 Create a virtual circuit.

```
mpls l2-vpn create dynamic-vc <dynamic-vc> pw-id <NUMBER:
1-2147483647> peer <IP address> tunnel <MPLS ingress
primary tunnel> pw-type <eth-raw|eth-tagged> pw-mode
<mesh|spoke>
```

where

dynamic-vc <dynamic-vc>	is the dynamic VS name.
pw-id<NUMBER: 1-2147483647>	is the VPN ID.
peer <IP address>	is the destination IP address.
tunnel <MPLS ingress primary tunnel>	is the transport tunnel.
pw-type <eth-raw|eth-tagged>	is the pseudowire type.
pw-mode <mesh|spoke>	is the pseudowire mode.

6 Attach the virtual circuit to the virtual switch.

```
virtual-switch ethernet attach mpls-vc <mpls-vc> vs <vs>
```

where

mpls-vc <mpls-vc>	is the list of MPLS VCs.
vs <vs>	is the virtual switch name.

—end—

Copyright© 2012-2013 Ciena® Corporation

Example

The following sample configuration file segment shows the VPWS configuration for the 3930 named R1 in the network topology diagram shown in Figure 11-9.

```
!------------------------------ vpws-r1.r4--------------------------!
! ----The primary path is created towards R2. This is also known as an Explicit
Route Object and is used to override the behavior of the IGP.
rsvp-te path create rsvp-path vpws-r1.r4.pri
rsvp-te path set rsvp-path vpws-r1.r4.pri index 5 ip 10.101.12.2 hop-type
loose

! ----Backup Path
rsvp-te path create rsvp-path vpws-r1.r4.bak
rsvp-te path set rsvp-path vpws-r1.r4.bak index 5 ip 10.104.13.3 hop-type
loose

! ----The primary tunnel is created as a dynamic RSVP tunnel. The destination
IP address is targeted at the loopback interface to ensure it is always up.
mpls tunnel create rsvp-ingress tun-vpws-r1.r4.pri dest-ip 4.4.4.4 path vpws-
r1.r4.pri record-route on
! ----The backup tunnel is associated with the primary tunnel and is only
forwarding if the primary tunnel fails.
mpls tunnel create rsvp-ingress tun-vpws-r1.r4.bak dest-ip 4.4.4.4 path vpws-
r1.r4.bak backup tun-vpws-r1.r4.pri

! ---- Virtual Circuit, Virtual Switch, and Attachment Circuit
mpls l2-vpn create dynamic-vc vpws-r1.r4 pw-id 100 peer 4.4.4.4 tunnel tun-
vpws-r1.r4.pri pw-mode mesh pw-type eth-raw

! ---- The Virtual Switch is configured in VPWS mode to prevent more than 1
attachment circuit from being added. This adheres to the strict definition of
VPWS only having a single AC.
virtual-switch ethernet create vs vpws-r1.r4 mode vpws
virtual-switch ethernet attach mpls-vc vpws-r1.r4 vs vpws-r1.r4
virtual-switch ethernet add vs vpws-r1.r4 port 10 vlan 100
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 11-52
VPLS configuration example

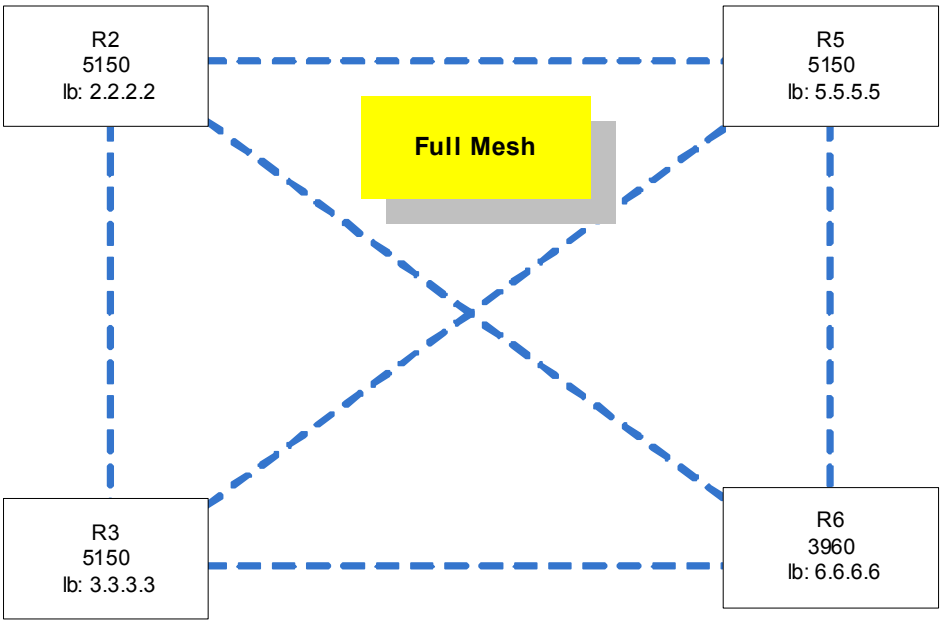
VPLS provides point-to-multipoint inter-LAN connectivity. With VPLS, PEs are connected to each other with a full mesh of virtual circuits for each VPLS instance. Each PE connects to an MTU-s on the UNI side through an MPLS network, or it connects to a CE device through an untagged interface or 802.1Q interface.

Note: In Release 6.9, the system handles the outer VLAN which could be C-VLAN, S-VLAN or B-VLAN and not the combination of multiple VLANs as in Q-in-Q or MAC-in-MAC. This is true for a UNI port on a PE also.

Each mesh virtual circuit connects to remote PE that is member of the same VPLS instance which provides emulated LAN service for the broadcast domain of the virtual switch. When configuring the mesh virtual circuit, include the virtual circuit identifier. For VPLS configuration, all the mesh virtual circuits that are associated with the same virtual switch, may or may not have the same value.

Note: In order to be interoperable with other vendors such as Cisco or Juniper, configure the same virtual circuit identifier for the mesh virtual circuits that are added to the same virtual circuit.

Figure 11-10
VPLS topology



Copyright© 2012-2013 Ciena® Corporation

To configure the VPLS service:

Step	Action
1	Create a virtual switch. `virtual-switch ethernet create vs <vs-name> mode <vpws|vpls>`
2	Add the UNI port(s) to the virtual switch. `virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} [statistics <on | off>] [translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100 | 9100 | 88A8>}`

3	Create an explicit path for each PE in the VPLS Mesh.
	`rsvp-te path create rsvp-path <name>`
	`rsvp-te path set rsvp-path <name> index <#> ip <ip> hop-type <loose|strict>`

4	Create an MPLS tunnel for each PE in the MPLS mesh. `mpls tunnel create rsvp-ingress <name> dest-ip <ip> path <name> record-route <on|off>`
5	Create a virtual circuit (Pseudowire) for each PE in the MPLS Mesh. `mpls l2-vpn create dynamic-vc <name> pw-id <#> peer <ip> tunnel <name> pw-mode <spoke|mesh> pw-type <eth-raw|eth-tagged>`
6	Attach each virtual circuit created in Step 5 to the virtual switch created in Step 1. `virtual-switch ethernet attach mpls-vc <name> vs <name>`

—end—

Example

The following sample configuration file segment shows the VPLS configuration for the 5150 named R3 shown in Figure 11-9.

```
                !----------------------vpls-100
virtual-switch ethernet create vs vpls-100 mode vpls
virtual-switch ethernet add vs vpls-100 port 8 vlan 100

                !---------------------------Peer: 2.2.2.2
                ! ----The path is associated with the tunnel. The purpose of the path is to allow
                configuration or strict routing for the LSP.
rsvp-te path create rsvp-path R3_R2.path
rsvp-te path set rsvp-path R3_R2.path index 5 ip 10.103.34.4 hop-type loose

                ! ----Record route information is used to calculate Fast Re-route paths and
                therefore should always be enabled if FRR protection is desired.
mpls tunnel create rsvp-ingress R3_R2 dest-ip 2.2.2.2 path R3_R2.path record-
route on

                ! ---- The pw-mode MESH is what allows the ELAN behavior associated with a
                VPLS. The pw-type ETH-RAW along with the attachment circuit type defines
                the treatment of the outer VLAN of the ingress frame.
mpls l2-vpn create dynamic-vc vpls-100.R2 pw-id 100 peer 2.2.2.2 tunnel R3_R2
pw-mode mesh pw-type eth-raw

virtual-switch ethernet attach mpls-vc vpls-100.R2 vs vpls-100

                !-------------------------Peer: 5.5.5.5
                ! ----Path to peer
rsvp-te path create rsvp-path R3_R5.path
rsvp-te path set rsvp-path R3_R5.path index 5 ip 10.103.34.4 hop-type loose

                ! ----MPLS Tunnel
mpls tunnel create rsvp-ingress R3_R5 dest-ip 5.5.5.5 path R3_R5.path record-
route on

                ! ---- Virtual Circuit
mpls l2-vpn create dynamic-vc vpls-100.R5 pw-id 100 peer 5.5.5.5 tunnel R3_R5
pw-mode mesh pw-type eth-raw

virtual-switch ethernet attach mpls-vc vpls-100.R5 vs vpls-100

                !-------------------------Peer: 6.6.6.6
                ! ----Path to peer
rsvp-te path create rsvp-path R3_R6.path
rsvp-te path set rsvp-path R3_R6.path index 5 ip 10.103.34.4 hop-type loose

                ! ----MPLS Tunnel
mpls tunnel create rsvp-ingress R3_R6 dest-ip 6.6.6.6 path R3_R6.path record-
route on

                ! ---- Virtual Circuit
mpls l2-vpn create dynamic-vc vpls-100.R6 pw-id 100 peer 6.6.6.6 tunnel R3_R6
pw-mode mesh pw-type eth-raw

virtual-switch ethernet attach mpls-vc vpls-100.R6 vs vpls-100
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 11-53
H-VPLS Configuration

A Hierarchical VPLS (H-VPLS) model is used to allow spoke connections to the VPLS core. A device provides the functionality to interface with the VPLS core by functioning as an MTU-s, which is connected as a spoke in the VPLS core using a virtual circuit.

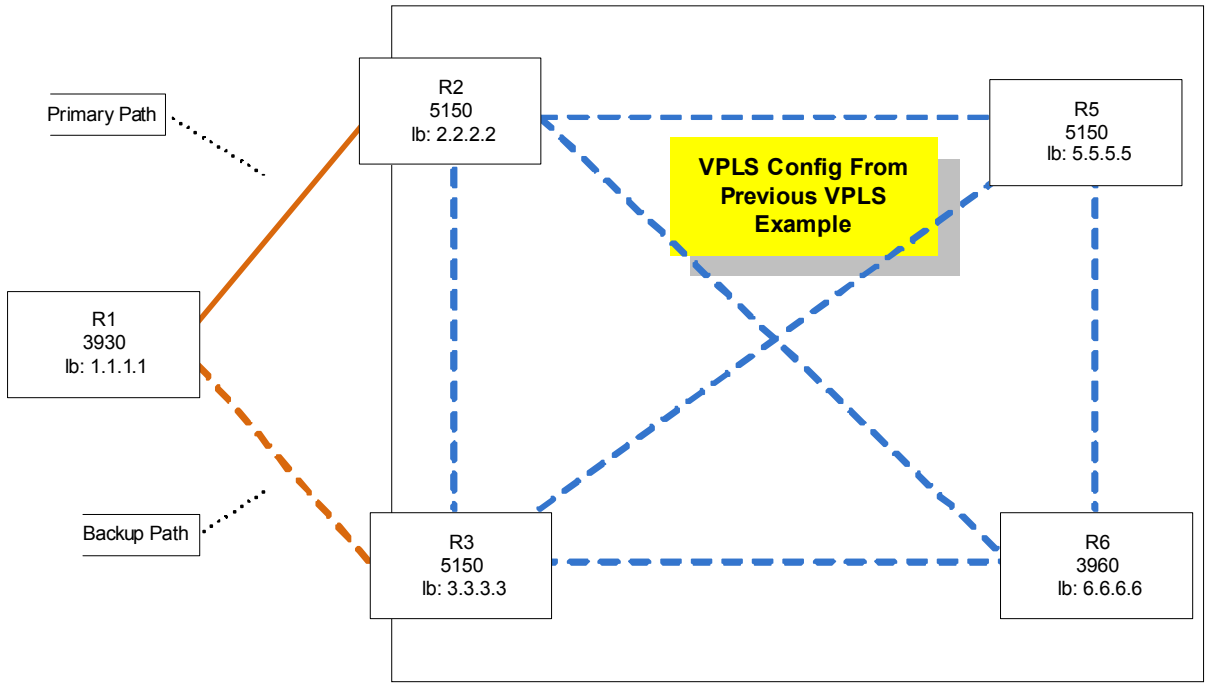
An MTU-s connects to an PE on the NNI side through an MPLS network and to a CE on the UNI side through an untagged interface or an 802.1Q interface.

> ***Note:*** In Release 6.9, the system handles the outer VLAN which could be C-VLAN, S-VLAN or B-VLAN and not the combination of multiple VLANs as in Q-in-Q or MAC-in-MAC. This is true for a UNI port on a PE also.

Although MTU-s is an external device to a PE, it is viewed as an adjunct device to the PE and used to address scalability.

Figure 11-11 illustrates the H-VPLS topology which includes VC redundancy for protection switching.

Figure 11-11
H-VPLS Topology



Copyright© 2012-2013 Ciena® Corporation

To configure the H-VPLS service:

Step	Action
1	Create a virtual switch. `virtual-switch ethernet create vs <vs-name> mode <vpws|vpls>`
2	Attach the UNI port(s) to the virtual switch. `virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-dot1dpri <NUMBER: 0-7>} [statistics <on | off>] [translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100 | 9100 | 88A8>}`

3	Create an explicit path for each endpoint.
	`rsvp-te path create rsvp-path <name>`
	`rsvp-te path set rsvp-path <name> index <#> ip <ip> hop-type <loose|strict>`

4	Create an MPLS tunnel for each endpoint. `mpls tunnel create rsvp-ingress <name> dest-ip <ip> path <name> record-route <on|off>`
5	Create a primary and protection virtual circuit (Pseudowire). `mpls l2-vpn create dynamic-vc <vc-name> pw-id <value> peer <peer-ip-addr> tunnel <tunnel-name>`
6	Attach the each virtual circuit to the virtual switch created in Step 1. `virtual-switch ethernet attach vs <vs-name> mpls-vc <vc-name>`

—end—

Example

The following sample configuration file segment shows the VPLS configuration for the 3930 named R1 shown in Figure 11-11.

```
                H-VPLS Configuration Example on R1
                !---------------hvpls-200
                !----Since VPLS is a multipoint service the VS mode should be set to VPLS
                which will allow multiple attachment circuits. VPWS mode only allows a single
                AC.
virtual-switch ethernet create vs hvpls-200 mode vpls
virtual-switch ethernet add vs hvpls-200 port 8 vlan 200

                !---------------Primary Virtual Circuit: hvpls-200P
                ! ---- Primary Path
rsvp-te path create rsvp-path R1_R2.path
rsvp-te path set rsvp-path R1_R2.path index 5 ip 10.101.12.2 hop-type loose

                !-----Primary Tunnel
mpls tunnel create rsvp-ingress R1_R2 dest-ip 2.2.2.2 path R1_R2.path record-
route on


                ! ---- Virtual Circuit, Virtual Switch, and Attachment Circuit
mpls l2-vpn create dynamic-vc hvpls-200P pw-id 200 peer 2.2.2.2 tunnel R1_R2
pw-mode spoke pw-type eth-raw

virtual-switch ethernet attach mpls-vc hvpls-200P vs hvpls-200

                ! ---- Backup Path
rsvp-te path create rsvp-path R1_R3.path
rsvp-te path set rsvp-path R1_R3.path index 5 ip 10.104.13.3 hop-type loose

                !-------Backup Tunnel
mpls tunnel create rsvp-ingress R1_R3 dest-ip 3.3.3.3 path R1_R3.path record-
route on

                ! ---- Virtual Circuit, Virtual Switch, and Attachment Circuit
mpls l2-vpn protection create dynamic-vc hvpls-200B primary-vc-name hvpls-
200P secondary-pw-id 200 peer 3.3.3.3 tunnel R1_R3
virtual-switch ethernet attach mpls-vc hvpls-200B vs hvpls-200
```


Copyright© 2012-2013 Ciena® Corporation

L2 Control Frame Tunneling configuration

This chapter explains how to configure the treatment of Layer 2 (L2) control frames with L2 control frame tunneling. With L2 Control Frames Tunneling, you can change the handling of untagged L2 control frames to be processed and forwarded as if they were data frames instead of being discarded or locally processed. Also, you can change the handling of transparent L2 control frames to be transformed to L2PT frame format.

> ***Note:*** To configure L2 control frame tunneling, you need to install the Advanced Ethernet license key. To obtain the Advanced Ethernet license key, contact Ciena Sales.

Overview

In conformance with IEEE standards, the default configuration does not propagate untagged L2 control frames through the network. If the default disposition is “discard”, the frames are dropped without further processing. If

Copyright© 2012-2013 Ciena® Corporation

the disposition is “peer”, the frames are forwarded to the software process that handles the specific protocol. Table 12-1 shows the frame format for the three forms for each protocol along with the default disposition for each.

Table 12-1
Control frame formats and disposition

Control frame name (with CLI attribute)	Form	Default disposition	DA MAC address	Type/Length encapsulation	LLC or EtherType	Subtype
Rapid Spanning Tree Protocol (rstp)	Untagged	Peer	01-80-C2-00-00-00	N/A	DSAP/ SSAP=0x42 Ctrl=0x03	N/A
	Tagged	Same as data frame	01-80-C2-00-00-00	N/A	DSAP/ SSAP=0x42 Ctrl=0x03	N/A
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	N/A	DSAP/ SSAP=0x42 Ctrl=0x03	N/A
Link Aggregation Control Protocol (lacp)	Untagged	Peer	01-80-C2-00-00-02	N/A	88-09	01
	Tagged	Same as data frame	01-80-C2-00-00-02	N/A	88-09	01
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	N/A	88-09	01
LACP Marker (lacp-marker)	Untagged	Peer	01-80-C2-00-00-02	N/A	88-09	02
	Tagged	Same as data frame	01-80-C2-00-00-02	N/A	88-09	02
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	N/A	88-09	02
802.3ah Operation, Administration, and Maintenance (oam)	Untagged	Discard	01-80-C2-00-00-02	N/A	88-09	03
	Tagged	Same as data frame	01-80-C2-00-00-02	N/A	88-09	03
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	N/A	88-09	03

Copyright© 2012-2013 Ciena® Corporation

Table 12-1
Control frame formats and disposition

Control frame name (with CLI attribute)	Form	Default disposition	DA MAC address	Type/Length encapsulation	LLC or EtherType	Subtype
IEEE Std. 802.1X/ EAPOL(802.1x)	Untagged	Discard	01-80-C2-00-00-03	N/A	88-8E	N/A
	Tagged	Same as data frame	01-80-C2-00-00-03	N/A	88-8E	N/A
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	N/A	88-8E	N/A
LLDP 802.1ab(lldp)	Untagged	Peer	01-80-C2-00-00-0E	N/A	88-CC	N/A
	Tagged	Same as data frame	01-80-C2-00-00-0E	N/A	88-CC	N/A
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	N/A	88-CC	N/A
Generic Attribute Registration Protocol (GARP) Block (garp-block)	Untagged	Discard	01-80-C2-00-00-2X	N/A	N/A	N/A
	Tagged	Same as data frame	01-80-C2-00-00-2X	N/A	N/A	N/A
	L2PT	Not supported	Not Supported	N/A	N/A	N/A
GARP Mulicast Registration Protocol (gmrp)	Untagged	Discard	01-80-C2-00-00-20	N/A	00-01	N/A
	Tagged	Same as data frame	01-80-C2-00-00-20	N/A	00-01	N/A
	L2PT	Not supported	Not Supported	N/A	N/A	N/A
GARP VLAN Registration Protocol (gvrp)	Untagged	Discard	01-80-C2-00-00-21	N/A	00-01	N/A
	Tagged	Same as data frame	01-80-C2-00-00-21	N/A	00-01	N/A
	L2PT	Not supported	Not Supported	N/A	N/A	N/A

Copyright© 2012-2013 Ciena® Corporation

Table 12-1
Control frame formats and disposition

Control frame name (with CLI attribute)	Form	Default disposition	DA MAC address	Type/Length encapsulation	LLC or EtherType	Subtype
Cisco-Port Aggregation Protocol (cisco-pagp)	Untagged	Discard	01-00-0C-CC-CC-CC	SNAP	01-04	N/A
	Tagged	Same as data frame	01-00-0C-CC-CC-CC	SNAP	01-04	N/A
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	SNAP	01-04	N/A
Cisco-Unidirectional Link Detection (cisco-udld)	Untagged	Discard	01-00-0C-CC-CC-CC	SNAP	01-11	N/A
	Tagged	Same as data frame	01-00-0C-CC-CC-CC	SNAP	01-11	N/A
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	SNAP	01-11	N/A
Cisco Discovery Protocol (cisco-cdp)	Untagged	Discard	01-00-0C-CC-CC-CC	SNAP	20-00	N/A
	Tagged	Same as data frame	01-00-0C-CC-CC-CC	SNAP	20-00	N/A
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	SNAP	20-00	N/A
Cisco-Dynamic Trunking Protocol (cisco-dtp)	Untagged	Discard	01-00-0C-CC-CC-CC	SNAP	20-04	N/A
	Tagged	Same as data frame	01-00-0C-CC-CC-CC	SNAP	20-04	N/A
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	SNAP	20-04	N/A
Cisco-VLAN Trunking Protocol (cisco-vtp)	Untagged	Discard	01-00-0C-CC-CC-CC	SNAP	20-03	N/A
	Tagged	Same as data frame	01-00-0C-CC-CC-CC	SNAP	20-03	N/A
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	SNAP	20-03	N/A

Copyright© 2012-2013 Ciena® Corporation

Table 12-1
Control frame formats and disposition

Control frame name (with CLI attribute)	Form	Default disposition	DA MAC address	Type/Length encapsulation	LLC or EtherType	Subtype
Cisco-Per VLAN Spanning Tree (cisco-pvst)	Untagged	Discard	01-00-0C-CC-CC-CD	SNAP	01-0B	N/A
	Tagged	Same as data frame	01-00-0C-CC-CC-CD	SNAP	01-0B	N/A
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	SNAP	01-0B	N/A
Cisco-VLAN Bridge (vlan-bridge)	Untagged	Discard	01-00-0C-CD-CD-CE	SNAP	01-0C	N/A
	Tagged	Same as data frame	01-00-0C-CD-CD-CE	SNAP	01-0C	N/A
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	SNAP	01-0C	N/A
Cisco STP Fast Uplink (cisco-stp-uplink-fast)	Untagged	Discard	01-00-0C-CD-CD-CD	SNAP	20-0A	N/A
	Tagged	Same as data frame	01-00-0C-CD-CD-CD	SNAP	20-0A	N/A
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	SNAP	20-0A	N/A
Bridge Block (bridge-block)	Untagged	Discard	01-80-C2-00-00-0X	N/A	N/A	N/A
	Tagged	Same as data frame	01-80-C2-00-00-0X	N/A	N/A	N/A
	L2PT	Not supported	Not Supported	N/A	N/A	N/A
All Bridges Block (all-bridges-block)	Untagged	Discard	01-80-C2-00-00-1X	N/A	N/A	N/A
	Tagged	Same as data frame	01-80-C2-00-00-1X	N/A	N/A	N/A
	L2PT	Not supported	Not Supported	N/A	N/A	N/A

Copyright© 2012-2013 Ciena® Corporation

Tunnel method

The tunnel method controls the classification and processing of untagged L2 control frames. You can set the tunnel method for a virtual switch to one of two modes, transparent or L2PT. The default mode is L2PT.

> ***Note:*** The maximum number of L2 control frame tunnels in a device is equal to the number of physical ports available, that is, UNI plus NNI.

Transparent Mode

When the tunnel method is in transparent mode, the switching chip hardware handles dispositions without sending frames to the CPU for processing, and is referred to as "fast path" operation. With "fast path" operation, untagged L2 control frame dispositions are maintained when the control frame ingresses a customer-facing interface with the untagged-ctrl-vs specified. Protocol dispositions are not maintained when the equivalent tagged L2 control frame ingresses a network-facing interface(s). In transparent mode, the L2 control frame is encapsulated accordingly, and the IEEE or Cisco MAC DA is left intact.

Transparent mode is only supported on MPLS EPL virtual switches. MPLS virtual switches do not support the L2PT mode.

L2PT Mode

When the virtual switch is configured for L2PT mode, L2 control frames are classified and either dropped or forwarded, according to each protocol's disposition and the untagged-ctrl-vs designation at each customer-facing interface. In L2PT mode the L2 control frame is fully decoded based on the proper IEEE or Cisco MAC DA and any relevant Op-Codes in the frame. If there is a specific protocol match on a protocol that supports L2PT stamping, the frame is eligible for conversion to and from L2PT and proper IEEE or Cisco MAC DA. Certain protocols do not support L2PT transformation, since their classification data does not contain enough information to uniquely identify the protocol once the protocol-specific MAC DA has been replaced. Such protocols do not have an L2PT form.

In general, when the virtual switch is in L2PT mode, the switching chip hardware sends L2 control frames to the CPU for "slow-path" processing. The frame agent will egress that frame out the customer-facing interfaces as a "proper" IEEE or Cisco L2 control frame, and egress the frame out network-facing interfaces as the same L2PT frame.

L2 control frames that ingress a network-facing interface are processed as follows:

- If the L2 control frame is in the L2PT-equivalent form and is fully decodable, it is converted back to the proper IEEE or Cisco MAC DA.

Copyright© 2012-2013 Ciena® Corporation

  - — Egressing Customer-Facing Interface(s): The proper IEEE or Cisco MAC DA form of the frame will egress out all customer-facing interfaces of the virtual switch.
  - — Egressing Network-Facing Interface(s): The L2PT form of the frame will egress out all network-facing interfaces, except for the interface on which the frame ingressed.
- If the L2 control frame has an L2PT MAC DA, but is not decodable, that frame will be switched out all customer-facing and network-facing interfaces with no conversion of MAC DA.
- If the L2 control frame is already in the proper IEEE or Cisco MAC DA form, that frame will be switched out all customer-facing and network-facing interfaces with no conversion of MAC DA to L2PT. This means that if an IEEE L2 control frame enters a network-facing interface, that same frame will egress interfaces regardless of the L2PT mode enabled on the virtual switch.

L2 control frames that ingress a customer-facing Interface are processed as follows:

- If the L2 control frame is in the L2PT-equivalent form, it will be treated as a data frame and switched out all customer-facing and network-facing interfaces with no conversion of MAC DA.
- If the L2 control frame is in the proper IEEE or Cisco MAC DA untagged form, that frame will subject to the disposition set for the protocol at the virtual switch L2 control frame tunneling instance. This untagged L2 control frame has classified to the customer facing port based upon the untagged-ctrl-vs attribute of that port.
  - — Egressing Customer-Facing Interface(s): The proper IEEE or Cisco MAC DA form of the frame will egress out all customer-facing interfaces of the virtual switch. From customer-to-customer, the L2PT form will not propagate.
  - — Egressing Network-Facing Interface(s): The L2PT form of the frame will egress out all network-facing interfaces, except for the interface on which the frame ingressed.
- If the L2 control frame has an IEEE or Cisco MAC DA, but is not decodable, that frame will be switched out all customer-facing and network-facing interfaces with no conversion of MAC DA.

Copyright© 2012-2013 Ciena® Corporation

Configuration

L2 control frame tunneling directs untagged L2 control frames into a forwarding domain of a virtual switch. Every virtual switch is associated with an L2 control frame tunneling configuration container. The container storing the configuration information is called an L2 control frame tunneling instance, which includes:

- **Administrative Status** - whether or not L2 control frame tunneling is enabled or disabled for the virtual switch.
- **Fixed .1D priority** - sets the .1D priority value applied to untagged L2 control frames to provide the baseline for Class of Service (CoS) treatment of the frame as it switches through the device. The priority value ranges from 0-7, where 7 has the highest priority and 0 has the lowest priority.
- **Protocol Disposition List** - sets a user-defined list of protocols to be processed by L2 control frame tunneling along with a disposition action of forward or discard.

For the related procedure, see “Configuring L2 Control Frame Tunneling” on page 12-19.

Copyright© 2012-2013 Ciena® Corporation

Procedure 12-1
Adding and removing untagged L2 control frame classification

In order for untagged L2 control frames to be directed from a port to a forwarding domain, you need to add the untagged L2 control frame traffic classification (untagged-ctrl-vs) to the port and the port must have some membership in the designated virtual switch `(virtual-switch ethernet add vs <VirtualSwitchName> port <PortName> vlan <VlanId>)`. When the untagged-ctrl-vs attribute has been designated, then the L2 control frame tunneling instance of the specified virtual switch will be used to handle the untagged L2 control frames that ingress the logical-port.

When the untagged-ctrl-vs attribute is removed, an untagged L2 control frame is processed at the port according the default disposition for its protocol.

Step	Action

To add L2 control frame classification

1 Add L2 control frame traffic classification for a port:

```
port set port <PortName> untagged-ctrl-vs
<VirtualSwitchName>
```

To remove L2 control frame traffic classification

2 Remove L2 control frame traffic classification for a port:

```
port unset port <PortName> untagged-ctrl-vs
<VirtualSwitchName>
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 12-2
Enabling and disabling L2 control frame tunneling

You can enable an L2 control frame tunnel on each port of each device simultaneously.

> ***Note:*** The maximum number of L2 control frame tunnels in a device is equal to the number of physical ports available, that is, UNI plus NNI.

When you disable L2 control frame tunneling for a specified virtual switch, any configured L2PT transforms will not occur, and untagged L2 control frames will be handled according to their default disposition.

Step	Action

To enable L2 control frame tunneling

1 Enable L2 control frame tunneling:

```
virtual-switch l2-cft enable vs <VirtualSwitchName>
```

where

vs <VirtualSwitchName>	is the virtual switch that you want to enable L2 control frame tunneling on

To disable L2 control frame tunneling

2 Disable L2 control frame tunneling:

```
virtual-switch l2-cft disable vs <VirtualSwitchName>
```

where

vs <VirtualSwitchName>	is the virtual switch that you want to disable L2 control frame tunneling on

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 12-3
Adding and removing control protocols

When adding a control protocol, you can optionally specify the disposition. If you do not specify the disposition, the default disposition for the control protocol is used.

When you remove a control protocol from a L2 control frame tunneling instance, untagged L2 control frames will be processed according to the default disposition for the control protocol.

Step	Action

To add a control protocol

1 Add a control protocol to an L2 control frame tunneling instance:

```
virtual-switch l2-cft protocol add vs <VirtualSwitchName>
{ctrl-protocol <802.1x|all-bridges-block|bridge-
block|cisco-cdp|cisco-dtp|cisco-pagp|cisco-pvst|cisco-
stp-uplink-fast|cisco-udld|cisco-vtp|garp-
block|gmrp|gvrp|lacp|lacp-markerlldp|oam|rstp|vlan-
bridge>} [disposition <discard|forward>]
```

where

vs <VirtualSwitchName>	is the virtual switch.
ctrl-protocol	is the control protocol.
disposition	Valid values are: • discard, which discards the frame without further processing • forward, which directs the frame to the forwarding domain associated with the L2 control frame tunneling instance. Frames with this disposition are subject to all ingress and egress transformations configured on the interface. In addition, these frames are flooded to all logical interfaces attached to the virtual switch associated with the L2 control frame tunneling instance.

To remove a control protocol

2 Remove a control protocol from an L2 control frame tunneling instance:

```
virtual-switch l2-cft protocol remove vs
<VirtualSwitchName> {ctrl-protocol <802.1x|all-bridges-
block|bridge-block|cisco-cdp|cisco-dtp|cisco-pagp|cisco-
```


Copyright© 2012-2013 Ciena® Corporation

```
pvst|cisco-stp-uplink-fast|cisco-udld|cisco-vtp|garp-
block|gmrp|gvrp|lacp|lacp-marker|lldp|oam|rstp|vlan-
bridge>}
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 12-4
Setting the disposition of control protocols

After you add a control protocol to an L2 control frame tunneling instance, you can modify its disposition.

Step	Action
1	Set the disposition of a control protocol in an L2 control frame tunneling instance: `virtual-switch l2-cft protocol set vs <VirtualSwitchName> {ctrl-protocol <802.1x|all-bridges-block|bridge-block|cisco-cdp|cisco-dtp|cisco-pagp|cisco-pvst|cisco-stp-uplink-fast|cisco-udld|cisco-vtp|garp-block|gmrp|gvrp|lacp|lacp-marker|lldp|oam|rstp|vlan-bridge>} {disposition <discard|forward>}` **—end—**

Copyright© 2012-2013 Ciena® Corporation

Procedure 12-5
Setting Layer 2 CFT attributes

When an untagged frame is forwarded (encapsulated) by the L2 control frame tunneling instance, it is given a Fixed .1D priority value. The default Fixed .1D priority value is 6. You can modify it to a value between 0 and 7.

The tunnel method controls the classification and processing of untagged L2 control frames. You can set the tunnel method for a virtual switch to one of two modes, transparent or L2PT. The default mode is L2PT.

Step	Action

1 Set Layer 2 CFT attributes:

```
virtual-switch l2-cft set vs <vs> priority <NUMBER: 0-7>
tunnel-method <l2pt|transparent>
```

where

vs <vs>	is the virtual switch name.
priority <NUMBER: 0-7>	is the .1D priority for encapsulated frames.
tunnel-method <l2pt|transparent>	tunnel encapsulation method.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 12-6
Displaying enabled L2 control frame tunneling instances

You can display a summary of all currently enabled L2 control frame tunneling instances or details about a single L2 control frame tunneling instance specified by virtual switch.

Note: L2 Control Frame Tunneling is also called Control Protocol Tunneling (CPT). The `virtual switch ethernet show vs <VirtualSwitchName>` command shows the status and tunnel method for L2 control frame tunneling as CPT Status and CPT Method.

Step	Action
1	To display a summary of enabled L2 control frame tunneling instances: `virtual-switch l2-cft show`
2	To display details about a specific enabled L2 control frame tunneling instance: `virtual-switch l2-cft show [vs <VirtualSwitchName>]`

—end—

Examples

The following example shows sample output for a summary of enabled L2 control frame tunneling instances.

```
> virtual-switch l2-cft show

+------------------------ L2 Ctrl Protocol Tunneling ---------------------+
| Virtual Switch   | Admin State | Tunnel Method                | Priority |
+------------------+-------------+------------------------------+----------+
| cft_0            | Disabled    | l2pt                         | 6        |
| cft_1            | Disabled    | l2pt                         | 6        |
| cft_2            | Disabled    | l2pt                         | 6        |
| cft-vs           | Disabled    | l2pt                         | 6        |
+------------------+-------------+------------------------------+----------+
```


Copyright© 2012-2013 Ciena® Corporation

The following example shows sample output for details about a specific enabled L2 control frame tunneling instance.

```
> virtual-switch l2-cft show vs cft-vs
+--------- L2 Ctrl Protocol Tunneling - cft-vs          -+
| Parameter              | Value                         |
+------------------------+-------------------------------+
| Admin State            | Enabled                       |
| Tunnel Method          | l2pt                          |
| Priority               | 6                             |
+------------------------+-------------------------------+
| Protocol Name          | Disposition                   |
+------------------------+-------------------------------+
| rstp                   | forward                       |
| lacp                   | forward                       |
+------------------------+-------------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 12-7
Displaying the L2 control frame classification for a port

Display the status of L2 control frame classification for a specific port.

Step	Action
1	Display the status of L2 control frame classification: `port show port <PortName>` where <PortName> is the port to display the status of L2 control frame classification for. The status of L2 control frame classification appears in the Untagged Ctrl VS field. **—end—**

Copyright© 2012-2013 Ciena® Corporation

Procedure 12-8
Displaying L2 control frame tunneling configuration in the configuration file

L2 control frame tunneling related configuration is saved to the configuration file in various sections, including:

- VIRTUAL-CIRCUIT CONFIG:
- VIRTUAL-SWITCH CONFIG:
- L2 CFT PROTOCOL CONFIG:

Step	Action
1	Display configuration information: `configuration show`

—end—

Example

The following example shows sample output for the configuration show command.

```
> configuration show
...
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
! VIRTUAL-CIRCUIT CONFIG:  virtual circuits
!
virtual-circuit ethernet create vc cft-vc0 vlan 100
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
! VIRTUAL-SWITCH CONFIG:
!
virtual-switch add reserved-vlan 1000
!
virtual-switch ethernet create vs cft-vs vc cft-vc0
!
virtual-switch ethernet add vs cft-vs port 9
!
port set port 9 untagged-ctrl-vs cft-vs
!
!

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
! L2 CFT CONFIG:
!
virtual-switch l2-cft enable vs cft-vs
!
virtual-switch l2-cft protocol add vs cft-vs ctrl-protocol rstp disposition
forward
virtual-switch l2-cft protocol add vs cft-vs ctrl-protocol lacp disposition
forward
!
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 12-9
Configuring L2 Control Frame Tunneling

In general, L2 control frame tunneling includes the following steps:

- Add a disposition of forward or discard for each protocol to be handled by L2CFT to the L2CFT Instance associated with the virtual switch to be used as a L2CF forwarding-domain.
- Enable the L2CFT Instance.
- Specify the untagged-ctrl-vs for each customer-facing virtual switch interface member. Those interface members will handle untagged L2CFs.
- Enable L2PT mode at the virtual switch if L2PT transforms are to occur for tunneled L2 control frames.
- Configure the L2 control frame fixed .1D priority value - this is the frame PCP that will go onto the frame when encapsulated.

  Note: For L2CFT over MPLS you do not need to specify the untagged-ctrl-vs for each customer-facing virtual switch interface member.

Step	Action
1	Create the VLAN. `vlan create vlan <VlanIdList> [name <String[31]>]` where <VlanIdList> — is the VLAN identifier. [name <String[31]>] — is the name of the VLAN. If no name is specified, a default name is assigned with the form of “VLAN#<VLANid>. For example if the VLAN ID is 1000 the VLAN name will be “VLAN1000”. ***Note:*** VLAN names cannot begin with an integer. For example, 3PUD is not valid, but PUD3 is allowed.
2	Add the ingress and egress ports to the VLAN. `vlan add vlan <VlanIdList> port <PortNameList>` where <VlanIdList> — is the VLAN created in step 1. <PortNameList> — is the list of ingress and egress ports.
3	To confirm the creation of the VLAN, enter the following command: `vlan show`

Copyright© 2012-2013 Ciena® Corporation

4 Confirm that the ports are members of the VLAN.

```
vlan show vlan <VlanIdList>
```

where

<VlanIdList> is the VLAN created in step 1.

5 Create a virtual circuit and assign it to the VLAN.

```
virtual-circuit ethernet create vc cft-vc vlan
<VlanIdList>
```

where

<VlanIdList> is the VLAN created in step 1.

6 To confirm the creation of the virtual circuit, enter the following command:

```
virtual-circuit show
```

7 To confirm that the VLAN is associated with the virtual circuit, enter the following command:

```
virtual-circuit show vc cft-vc
```

8 Add a reserved VLAN.

```
virtual-switch add reserved-vlan <VLAN IDList>
```

where

<VLAN IDList> is the reserved VLANs used to create virtual switch Instances.

Note: This command is platform-dependent.

9 Create the virtual switch.

```
virtual-switch ethernet create vs <VirtualSwitchName[15]>
vc <VirtualCircuitName[15]>
```

where

vs <VirtualSwitch Name[15]> is the name of the virtual switch.

vc <VirtualCircuit Name[15]> is the name of the virtual circuit to be associated with the virtual switch.

10 Add the ingress port to the virtual switch.

```
virtual-switch ethernet add vs <vs> {port <port>} {vlan
<VLAN list>} {ip-interface <interface-name>} {encap-cos-
policy <dot1dpri-inherit | fixed | ip-prec-inherit |
phbg-inherit | port-inherit | vs-inherit>} {encap-fixed-
dot1dpri <NUMBER: 0-7>} [statistics <on | off>]
[translate-tag <NUMBER: 0-4094>] {service-vlan-tpid <8100
| 9100 | 88A8>}
```

where

vs <vs> is the virtual switch created in step 9.

<PortNameList> is the port name

Copyright© 2012-2013 Ciena® Corporation

11 Enable L2 control frame tunneling for the virtual switch.

```
virtual-switch l2-cft enable vs <VirtualSwitchName[15]>
```

where

vs <VirtualSwitch Name[15]>	is the virtual switch created in step 9.

12 Add the control protocols to the protocol disposition list with the desired disposition for the L2 control frame tunneling instance.

```
virtual-switch l2-cft protocol add vs
<VirtualSwitchName[15]> {ctrl-protocol <802.1x|all-
bridgesblock|bridge-block|cisco-cdp|ciscodtp|cisco-
pagp|cisco-pvst|cisco-stpuplink- fast|cisco-udld|cisco-
vtp|garpblock|gmrp|gvrp|lacp|lacpmarker|
lldp|oam|rstp|vlan-bridge>} [disposition
<discard|forward>]
```

where

vs <VirtualSwitch Name[15]>	is the virtual switch created in step 9.
{ctrl-protocol <802.1x|all-bridgesblock| bridge-block|cisco-cdp|ciscodtp| cisco-pagp|cisco-pvst|cisco-stpuplink-fast|cisco-udld|cisco-vtp|garpblock| gmrp|gvrp|lacp|lacpmarker| lldp|oam|rstp|vlan-bridge>}	is the control protocol to add.
[disposition <discard|forward>]	is the disposition. If left unspecified, the disposition is set to the default disposition for the specified control protocol.

13 Add the L2 control frame classification to the port for ingress traffic.

```
port set port <PortNameList> untagged-ctrl-vs
<VirtualSwitchName[15]>
```

where

port <PortNameList>	is the port created in step 10.
untagged-ctrl-vs <VirtualSwitch Name[15]>	is the virtual switch created in step 9.

Copyright© 2012-2013 Ciena® Corporation

14 To confirm the L2 control frame tunnel configuration of the virtual switch, enter the following command:

```
virtual-switch l2-cft show vs <VirtualSwitchName[15]>
```

where

untagged-ctrl-vs <VirtualSwitch Name[15]>	is the virtual switch created in step 9.

—end—

Example

The following example shows how you can configure virtual switches and ports to classify untagged L2 control frames and then transform them into the L2PT form. The ingress port classification allows untagged L2 control frames and frames with VLAN tag 100. The L2 control frame tunneling instance is enabled for the virtual switch. Also, RSTP and LACP are added to the protocol disposition list with the disposition of forwarding.

Create the VLAN 100.

```
vlan create vlan 100
```

Add the ingress and egress ports to VLAN 100.

```
vlan add vlan 100 port 1
```

Confirm the creation of VLAN 100.

```
vlan show
```

Confirm the ports are members of VLAN 100.

```
vlan show vlan 100
```

Create a virtual circuit and assign it to VLAN 100.

```
virtual-circuit ethernet create vc cft-vc vlan 100
```

Confirm the creation of the virtual circuit.

```
virtual-circuit show
```

```
      +------ ETHERNET VIRTUAL CIRCUIT TABLE -----+
      | Name             | VLAN                   |
      +------------------+------------------------+
      | cft-vc           | 100                    |
      +------------------+------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Confirm that VLAN 100 is associated with the virtual circuit.

```
virtual-circuit show vc cft-vc
```

```
      +------------ ETHERNET VIRTUAL CIRCUIT INFO ----------+
      | Parameter            | Value                        |
      +----------------------+------------------------------+
      | Name                 | cft-vc                       |
      | ID                   | 1                            |
      | Provider VLAN        | (100)    VLAN#100            |
      | Psuedo-Wire Type     | untagged                     |
      | VS Using             |                              |
      +-----------------------------------------------------+
```

Add reserved VLAN 1000.

```
virtual-switch add reserved-vlan 1000
```

Create the virtual switch.

```
virtual-switch ethernet create vs cft-vs vc cft-vc
```

Add the ingress port to the virtual switch.

```
virtual-switch ethernet add vs cft-vs port 1
```

Enable L2 control frame tunneling for the virtual switch.

```
virtual-switch l2-cft enable vs cft-vs
```

Add the control protocols to the protocol disposition list with the desired disposition for the L2 control frame tunneling instance.

```
virtual-switch l2-cft protocol add vs cft-vs ctrl-
protocol rstp disposition forward
```

```
virtual-switch l2-cft protocol add vs cft-vs ctrl-
protocol lacp disposition forward
```

Add the L2 control frame classification to the port for ingress traffic.

```
port set port 1 untagged-ctrl-vs cft-vs
```

Confirm the L2 control frame tunnel configuration of the virtual switch.

Copyright© 2012-2013 Ciena® Corporation

```
virtual-switch l2-cft show vs cft-vs

      +--------- L2 Ctrl Protocol Tunneling - cft-vs       -+
      | Parameter              | Value                      |
      +------------------------+----------------------------+
      | Admin State            | Enabled                    |
      | Tunnel Method          | l2pt                       |
      | Priority               | 6                          |
      +------------------------+----------------------------+
      | Protocol Name          | Disposition                |
      +------------------------+----------------------------+
      | rstp                   | forward                    |
      | lacp                   | forward                    |
      +------------------------+----------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Quality of Service configuration

This chapter details Quality of Service (QoS) features of the system software, including:

- "Class of Services (CoS) policies and mapping"
- "Traffic profiling"
- "Traffic profiling with hierarchical ingress metering"
- "Congestion management"
- "Egress scheduling"
- "Egress shaping"

Copyright© 2012-2013 Ciena® Corporation

Class of Services (CoS) policies and mapping

The system software associates an internal resolved CoS (R-CoS) value and resolved Color (R-COLOR) with every ingress frame that classifies to a port. The R-CoS and R-COLOR values provide the baseline for CoS treatment of a frame as it switches through the device. The R-CoS value ranges from 0-7, where 7 has the highest priority and 0 has the lowest priority. R-COLOR values are green and yellow. The initial R-CoS and R-Color values depend upon the configuration of the of the port, including virtual switch memberships.

As of release 6.9.1, traffic profiles support the ability to remark the meter's colored output from yellow to green, green to yellow, or leave as-is on a per traffic profile basis. In addition, the RCOS value can be remarked based on a frame's meter color output. See "Traffic profiling" on page 4 for more details.

The resolved CoS Policy configured on a port is ignored when it is involved in one or more virtual switch (subscriber) memberships. With that configuration, the R-CoS and R-Color is determined by the encap CoS policy (and priority) that is configured at the virtual switch. However, when the encap CoS policy configured on a Virtual Switch Member is port-inherit, the resolved Cos Policy of the port is applied.

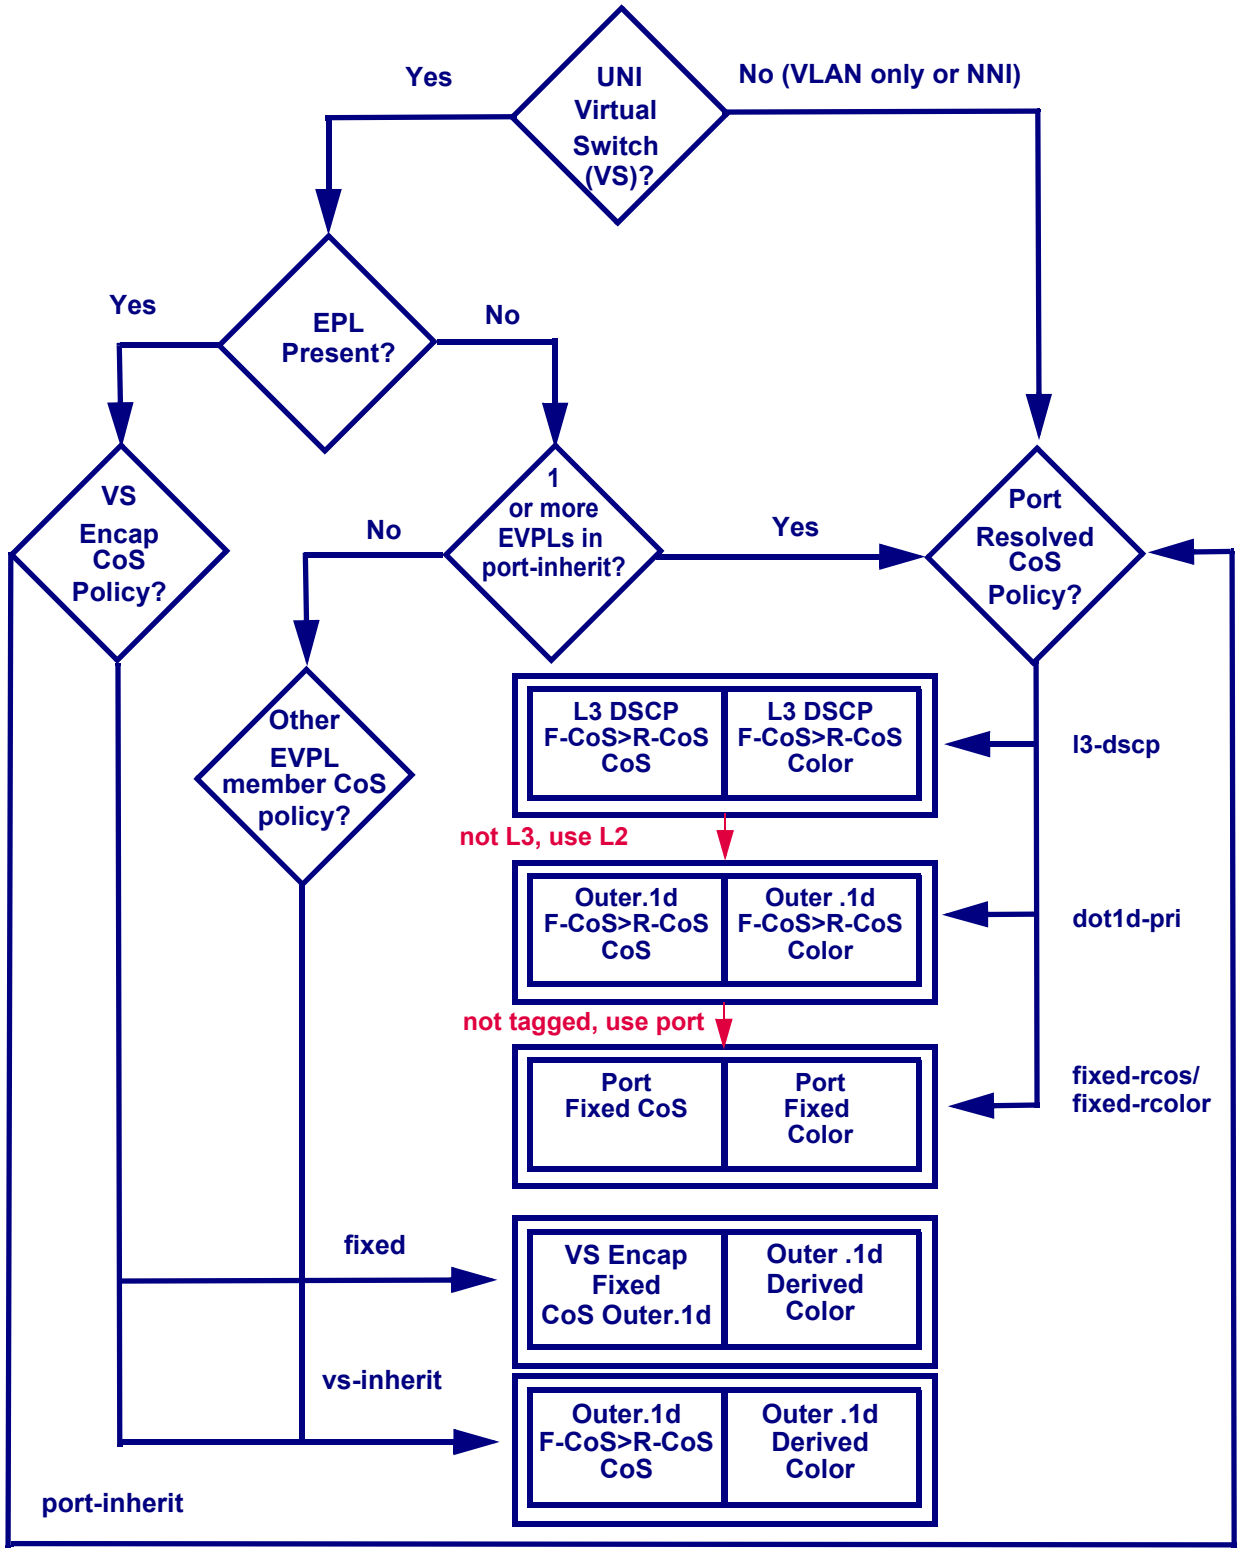
> ***Note:*** If an EVPL member is set to use the port-inherit policy is set to l3-dscp, then all of the L3 frames in the other EVPLs follow the L3 mapping table, even if set to vs-inherit or fixed. If the frames are not L3 and are tagged, then the L2 portion of the mapping table is in effect. If the frames are not L3 and are untagged without an IP header, then the R-CoS and R-Color are set according to the port's fixed setting. For untagged frames, the R-CoS and R-Color are set according to the port's fixed setting.

The Egress Frame CoS Policy is configured per port and defines when to use the Egress Frame CoS Map to remark an egress L2 frame's PCP and DEI fields based on the internal R-CoS and R-Color values.

Figure 13-1 shows which policy and portion of the map is in effect depending on the port's configuration.

Copyright© 2012-2013 Ciena® Corporation

Figure 13-1 Resolved CoS Policy and Mapping



Traffic profiling

Traffic profiling classifies traffic based on different classification rules and meters the traffic flow to configured EIR/PIR and CIR values defined in kbps. Depending upon how the traffic flow compares to these rates, the R-COLOR of each frame is set to a specific value upon ingress. Traffic ingressing at a rate:

- Up to CIR will be marked Green and allowed through.
- Above CIR and less than PIR will be marked Yellow and allowed through.
- Above PIR will be marked Red and dropped.

Note: On a 100 Mbps port the CIR cannot be set to the full amount of physical bandwidth of the port (100 Mbps). This is because 1 Mb of bandwidth is reserved for the ingress and egress of BPDUs and other high priority traffic. To commit all of the physical bandwidth to a CIR could create a serious condition in which the device does not receive critical frames. To avoid this issue, the CIR should only be configured to a maximum of 99 Mbps.

Determining classification attributes

In order for traffic to classify to a traffic profile, you need to configure classifiers for the selected mode by setting the port's standard profile. Even though classifiers are not mandatory upon creation, the traffic profile will not work without configuring appropriate classifiers. Classifiers include:

- dot1dpri - Classifies to the traffic profile based on 802.1D priority value ranging 0-7. Up to 8 classifiers are supported per port.
- ip-prec - Classifies to the traffic profile based on IP precedence value ranging 0-7. Up to 8 classifiers are supported per port.
- dscp - Classifies to the traffic profile based on DSCP. Up to 64 classifiers are supported per port.
- dscp-remark-policy - leave ingressing SVLAN 802.1p bit value as-is or mark with a fixed value.
- vlan - Classifies to the traffic profile based on VLAN ID. One VLAN classifier is supported per port per profile.
- vs - Classifies the traffic profile based on the Ethernet VS ID. Note that traffic profiles cannot classify to an MPLS VS.
- untagged - Classifies to the traffic profile if the traffic has no VLAN tag. Any other classifiers will be ignored, however, they will consume resources, if configured. To conserve resources when using the untagged classifier, clear any other classifiers from the traffic profile. One untagged classifier is supported per port per profile.

Copyright© 2012-2013 Ciena® Corporation

Note 1: Untagged traffic classifies first by the selected “mode.” Unless a profile is configured to specifically classify untagged traffic, untagged traffic with an IP header can classify to an IP precedence or DSCP based traffic profile,

Note 2: Traffic profiles with the untagged classifier cannot be assigned to ARP or non-conforming per port entries.

VLAN classifiers can be combined with 802.1D, IP precedence, or DSCP, depending upon the allocation of hardware resources. The maximum number of classifiers (including ARP and non-conforming) that are configurable per mode and per port (without any adjustment to resource management) are listed in Table 13-1.

Table 13-1
Configurable classifiers

Classifier	Per mode	Per port (3902, 3911, 3920, 3940, 5140)	Per port (3916, 3930, 3931)	Per port (3960, 5150)
dot1dpri	8	8	8	8
ip-prec (ipp)	8	8	8	8
dscp	64	64	64	64
vlan or vs*	1	64	1024	4096
vlan or vs/dot1dpri*	8	128	1024	4096
vlan or vs/ip-prec*	8	128	1024	4096
vlan or vs/dscp*	64	128	1024	4096
untagged	1	1	1	1

Note 1: The actual number of configurable classifiers will be less due to reserved profiles (ARP/non-conforming) per port. Each port automatically consumes 2 classifiers, 1 for ARP and 1 for non-conforming standard profiles.

Note 2: When setting multiple classifiers in one set command an overlap condition can occur, depending on the order in which the classifiers are added. To set both a CoS (dot1d, DSCP, IP) classifier with a VLAN classifier, set the CoS classifier first and then the VLAN classifier. For more information aboout traffic profiling error codes, refer to “Traffic profiling error codes”.

Copyright© 2012-2013 Ciena® Corporation

Note 3: *The VS maximum is up to the configurable limits of the system. The device will reach the VS limit before the classification limit is reached. For example, on a 3960 512 virtual switches can be configured even though 4096 classifiers can be configured.

Determining Metering Attributes

As of release 6.9.1, the meter provisioning mode can be set to either EIR or PIR, and PIR will begin to be deprecated. The mode that is chosen dictates how provisioning is accepted, displayed, and saved. If EIR is specified, PIR provisioning is disallowed.

Metering is also defined at the port's standard profle by the following attributes:

- cir - Committed information rate in kbits ranging from 0 up to the administrative port speed. (Mandatory, if PIR is not specified).
- eir - Excess information rate in kbits ranging from 0 up to the administrative port speed. (PIR = CIR+EIR when EIR is specified)
- pir - Peak information rate in kbits ranging from 0 up to the administrative port speed. (Mandatory, if CIR is not specified).

  Note: On 3902, 3911, 3940, and 5140 switches, the configured CIR and PIR always normalize to a rate equal to or greater than the configured value in 64 kbps increments in order to be processed by the hardware. Also, the PIR will normalize to be equal to or greater than the configured CIR. On 3916, 3930, 3931, 3960, and 5150 switches, the configured CIR and PIR always normalize to a rate in varying increments depending upon the sum of the configured rates of the CBS and EBS and the configured rate of the PIR.

  If the sum of the CBS and EBS is:

  — 0 or less than 16 M, the CIR and PIR increment in steps of 64 kbps

  — 16 M and less than 32 M, the CIR and PIR increment by 128 kbps

  — 32 M and less than 64 M, the CIR and PIR increment by 256 kbps

  — 64 M and less than 128 M, the CIR and PIR increment by 512kbps

  — 128 M and less than 512 M, the CIR and PIR increment by 1024 kbps

  If the PIR is:

  — 0 or less than 16.78 Gbps, the CIR and PIR increment in steps of 64 kbps

  — 16.78 Gbps and less than 33.55 Gbps, CIR and PIR increment by 128 kbps

  — 33.55 Gbps and less than 40 Gbps, CIR and PIR increment by 256 kbps

Copyright© 2012-2013 Ciena® Corporation

- cbs - Committed Burst Size (CBS) in kilobytes (kbytes). On 3902, 3911, 3940, and 5140 switches, the value entered for the CBS is adjusted by rounding up to the nearest power of 2 from 4 (2^2) up to 16384 (2^{14}) in Kbytes. On 3916, 3930, 3931, 3960, and 5150 switches, the CBS entered value, ranging from 0-262144 Kbytes, is adjusted by rounding up depending upon the sum of the CBS and EBS and the configured rate of the PIR using the same criteria as previously described for the CIR and PIR.
- ebs - Excess burst size (EBS) in kbytes. On 3902, 3911, 3940, and 5140 switches, the EBS entered is adjusted to fit the Peak Burst Sized (PBS), as defined in hardware based upon the CBS value, minus the CBS value. The PBS is rounded by rounding up to the nearest power of 2 from 4 (2^2) up to 16384 (2^{14}) in Kbytes, and the EBS is rounded to fit these increments when added to the CBS. For example, entering CBS of 4 Kbytes, and EBS of 8 Kbytes would equal a PBS of 12 Kbytes. Since the nearest power of 2 increment to 12 is 16, the EBS is rounded up to 12 Kbytes. On 3916, 3930, 3931, 3960, and 5150 switches, the EBS entered value, ranging from 0-262144 Kbytes, is the configured rate of the PIR using the same criteria as previously described for the CIR and PIR.

  Note: If CIR is equal to PIR, and CBS and EBS values are not specified, the system applies the default burst to both CBS and EBS upon a reboot. This does not apply to HIM profiles.

For the related procedures, see "Configuring per-port standard traffic profiling" on page 13-53 and "Configuring per-port and per-VLAN standard traffic profiling" on page 13-54.

Traffic profiling with hierarchical ingress metering

On the 3916, 3930, 3931, 3960, and 5150 platforms, traffic profiling supports HIM (hierarchical ingress metering). With hierarchical ingress metering, frames classify to a parent traffic profile for metering, and are fed to a child traffic profile for additional metering. Hierarchical ingress metering is supported by setting the per-port traffic profiling mode to one of the following types:

- Hierarchical port (hierarchical-port) - for port based classification at the parent traffic profile.

  Note: Untagged traffic is treated as non-conforming traffic and passes through the "default" profile at the child level, unless a child profile is configured to classify untagged traffic,

- Hierarchical VLAN (hierarchical-vlan) - for unique VLANs or a "default" classification for the parent traffic profile. Only one "default" profile may exist per parent or child. No other classification in the default profile.

Copyright© 2012-2013 Ciena® Corporation

- Advanced - supports both hierarchical and non-hierarchical traffic. See Advanced classification mode.

  Note 1: Untagged traffic is treated as non-conforming traffic and passes through the “default” profile at the parent level, unless a child profile or a separate profile at the parent level (above the default) is configured to classify untagged traffic.

Advanced classification mode

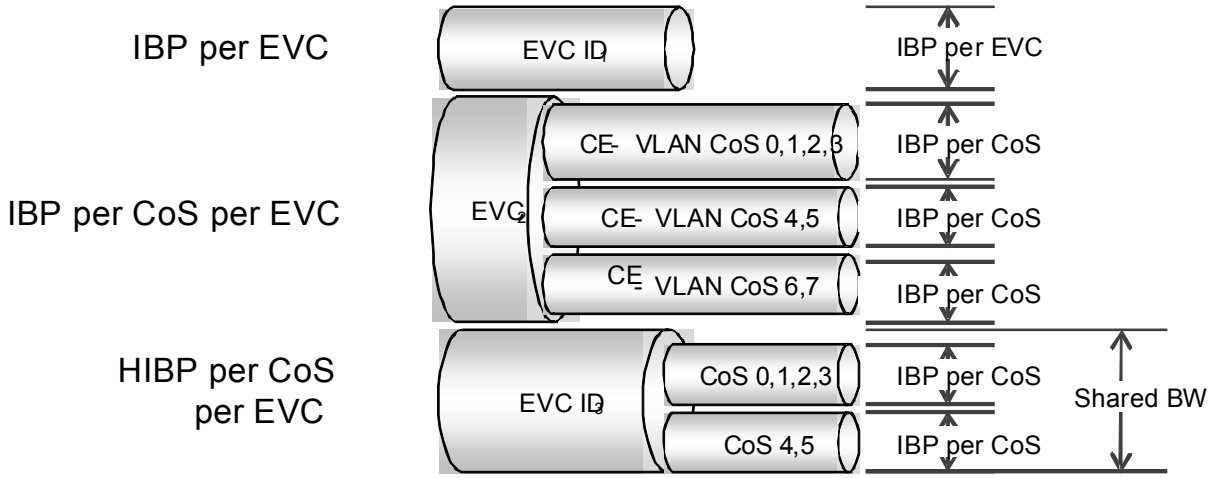
In the advanced traffic profile classification mode, both hierarchical and non-hierarchical traffic profiles can coexist. In addition to being able to configure HIM and standard-profiles on the same port, the system can classify on a VS name (or VS+VLAN, VS+COS and/or VS+VLAN+COS), as shown in Figure 13-2. Advanced mode is the only mode that supports this type of EVC based classification.

When configuring EVC-based traffic-profiling, the following dependencies must be considered:

- If the traffic-profile is classifying on a VS, the port that the profile is created on must be a member (EPL or EVPL) of the Virtual Switch that is being used as a classifier. Otherwise an error is generated when trying to create this traffic-profile.
- If the traffic-profile is classifying on VS+VLAN, the port that the profile is created on must be a member of the VS. If it is added as an EVPL, the VLAN used in the traffic-profile must match the VLAN of the VS member.
- VS members cannot be removed if the port that they are on still has traffic-profiles classifying on the VS. All traffic-profiles on the port must be deleted before removing the port as a member of the VS.
- A VS cannot be deleted if it still has traffic-profiles using it as a classifier. The traffic-profiles must be deleted first and then the VS can be deleted.

Copyright© 2012-2013 Ciena® Corporation

Figure 13-2
EVC advanced mode classification



When advanced mode is configured, the parent ports must have unique virtual switches across all parent profiles. There can be no overlap in parent classification. Also, child level classification cannot overlap any other classifier in use, that is, the profiles cannot overlap with other parent/child or other non-hierarchical profiles. A T2 untagged profile is the only exception as it could overlap other DSCP/IPP child classifiers under a given parent.

If a profile is created without a parent in advanced mode, it is assumed to be a standard profile without hierarchical metering, and it cannot be added to a parent later. It must be deleted and re-created with the parent associated at the time of creation. In addition, a child profile can only be removed from a parent by deleting the child profile.

> ***Note:*** In advanced mode if the parent is using a VLAN classifier, only h-vlan child modes are allowed, and vlan-cos generates an error. Likewise, if the parent classification is a VS, then any child mode other than vlan-cos generates an error.

For the related procedures, see "Configuring hierarchical VLAN traffic profiles" on page 13-56, "Configuring hierarchical port traffic profiles" on page 13-61, "Configuring VS classification for standard traffic profiles" on page 13-63, and "Configuring VS classification for HIM traffic profiles" on page 13-64.

With standard traffic profiling, a single port cannot support duplicate classification, such as the same VLAN or untagged. Hierarchical ingress metering provides the ability to classify based on multiple classifications and values to the same VLAN at different metering rates for the same port.

Copyright© 2012-2013 Ciena® Corporation

When changing to or from a hierarchical mode, all traffic profile must be removed from the port. Also, the arp-standard-profile and nonconform-standard-profile per port attributes do not apply to ports in a hierarchical traffic profiling mode. Before changing the parent's child mode, be sure to remove any child profiles.

Note: The classifier-mode attribute does not affect the "default" profile at the child level for hierarchical port mode. Also, it does not affect the "default" profile at the parent level for hierarchical VLAN mode.

Congestion management

Congestion management processing determines whether a frame should be enqueued or dropped. Congestion management is configured per queue based on a congestion avoidance profile associated with each CoS queue. This profile contains the drop parameters for traffic. Drop parameters are defined by the following general attributes:

- **Drop Threshold** - defines the percentage of the queue capacity that is reached before packets are eligible to be dropped.
- **Drop Probability** - define the percentage of packets that are dropped once the threshold is exceeded.

By dropping packets before the queue fills up, congestion management controls the average buffer queue size. As packets are dropped, the packet source automatically decreases its transmission rate until all the packets reach their destination, indicating that the congestion is cleared. TCP is an example of a protocol that dynamically adapts to congestion-induced frame loss.

Depending upon the hardware platform, the system software supports Simple Random Early Detection (sRED) profiles or simple Weighted Random Early Discard (sWRED) profiles for congestion avoidance that are enabled by default. You can configure custom congestion avoidance profiles for your platform and configure the egress port queues to use them.

Creating and modifying sRED profiles

The 3902, 3911, 3920, 3940, and 5140 platforms support sRED profiles. By default there is one congestion avoidance profile (Default-SRED) in the system that all queues reference, and each port supports two default sRED curves per queue, one for yellow traffic and one for green traffic. Each sRED curve supports a lower drop threshold and drop probability. The upper threshold is not configurable, and supports 128 packets per port.

When the number of packets are below the lower threshold, no packets are dropped. When they exceed the lower threshold, and are below the upper threshold, they are dropped according to the drop probability value. When the

Copyright© 2012-2013 Ciena® Corporation

number of packets reaches the upper threshold, all packets are dropped (100 percent drop probability) until the number of packets is below the upper threshold.

For yellow packets, the default drop threshold is 50 percent. For green packets, the default drop threshold is 75 percent. This means that the queue capacity can be filled with 75 percent green packets, and 50 percent yellow packets before the packets are eligible to be dropped. Additionally, both sRED curves assign a drop probability percentage for green (0.09765625 percent or 1/1024) and yellow (6.25 percent or 1/16) packets per CoS queue that is applied via a congestion avoidance profile.You can create up to 7 custom congestion avoidance profiles and modify their configuration.

Lower drop threshold values for green or yellow traffic range from 1-100 percent. Drop probability values for green or yellow traffic include:

- 100pct - 100 percent
- 6.25pct - 6.25 percent or 1/16.
- 3.125pct - 3.125 percent or 1/32.
- 1.5625pct - 1.5625 percent or 1/64.
- 0.78125pct - 0.78125 percent or 1/128.
- 0.390625pct - 0.390625 percent or 1/256.
- 0.1953125pct - 0.1953125 percent or 1/512.
- 0.09765625pct - 0.09765625 percent or 1/1024.

If you create a custom sRED profile without specifying the thresholds or drop probability values, the profile will be the same as the Default-SRED profile.

Creating and modifying sWRED profiles on the 3960

The 3960 platform supports simple Weighted Random Early Discard (sWRED) profiles to handle three types of traffic:

- TCP green traffic
- TCP yellow traffic (has been metered to yellow from traffic profiling)
- Non-TCP traffic

For each traffic type, an sWRED profile supports an sWRED curve with a configurable start threshold (1-100 percent), upper threshold (1-100 percent), and a maximum drop probability. Maximum drop probability values include:

- 100pct - 100 percent
- 75pct - 75 percent
- 50pct - 50 percent
- 25pct - 25 percent

Copyright© 2012-2013 Ciena® Corporation

- 10pct - 10 percent
- 9pct - 9 percent
- 8pct - 8 percent
- 7pct - 7 percent
- 6pct - 6 percent
- 5pct - 5 percent
- 4pct - 4 percent
- 3pct - 3 percent
- 2pct - 2 percent
- 1pct - 1 percent
- 0pct - 0 percent

When the number of packets are below the start threshold, no packets are dropped. When they exceed the start threshold, and are below the upper threshold, they are dropped according to the maximum drop probability value. When the number of packets reaches the upper threshold, all packets are dropped according to the maximum drop probability until the number of packets is below the upper threshold.

The default sWRED profile (Default-S-WRED) on the 3960 is configured as follows:

- TCP green start threshold - 75 percent
- TCP green maximum upper threshold - 100 percent
- TCP green maximum drop probability - 0 percent
- TCP yellow start threshold - 50 percent
- TCP yellow maximum upper threshold - 100 percent
- TCP yellow maximum drop probability - 100 percent
- Non-TCP start threshold - 50 percent
- Non-TCP maximum upper threshold - 100 percent
- Non-TCP maximum drop probability - 100 percent

If you create a custom sWRED profile without specifying the thresholds or drop probability values, the profile will be the same as the Default-S-WRED profile.

Creating and modifying sWRED profiles on the 3916, 3930, 3931, and 5150

The 3916, 3930, 3931, and 5150 platforms support sWRED profiles to handle two types of traffic:

- Green traffic

Copyright© 2012-2013 Ciena® Corporation

- Yellow traffic (has been metered to yellow from traffic profiling)

For each traffic type, an sWRED profile supports an sWRED curve with a configurable lower threshold (1-100 percent), upper threshold (1-100 percent), and a maximum drop probability. Maximum drop probability values include:

- 100pct - 100 percent
- 75pct - 75 percent
- 50pct - 50 percent
- 25pct - 25 percent
- 10pct - 10 percent
- 9pct - 9 percent
- 8pct - 8 percent
- 7pct - 7 percent
- 6pct - 6 percent
- 5pct - 5 percent
- 4pct - 4 percent
- 3pct - 3 percent
- 2pct - 2 percent
- 1pctv1 percent
- 0pct - 0 percent

When the number of packets are below the lower threshold, no packets are dropped. When they exceed the lower threshold, and are below the upper threshold, they are dropped according to the maximum drop probability value. When the number of packets reaches the upper threshold, all packets are dropped according to the maximum drop probability until the number of packets is below the upper threshold.

The default sWRED profile (Default-S-WRED) is configured as follows:

- Green lower threshold - 75 percent
- Green upper threshold - 100 percent
- Green maximum drop probability - 0 percent
- Yellow lower threshold - 50 percent
- Yellow upper threshold - 100 percent
- Yellow maximum drop probability - 100 percent

If you create a custom WRED profile without specifying the thresholds or drop probability values, the profile will be the same as the Default-S-WRED profile.

Copyright© 2012-2013 Ciena® Corporation

Egress scheduling

Scheduling determines the order in which the physical queues are processed. By default, the system software follows the strict priority scheduling. CoS queue 7 has the highest priority while CoS queue 0 has the lowest priority. When contention (or congestion) for bandwidth occurs in the system, the higher priority CoS queues will be serviced before the lower priority CoS queues. With the default scheduler algorithm (strict), the 802.1p priority (VLAN priority in the frame) is used to map to an internal priority (resolved CoS) which is then mapped to a physical CoS queue. The result is that frames with a higher priority will transmit before those with a lower priority. Strict priority queuing reduces resources for low priority packets, but ensures low-latency servicing of high-priority packets.

In addition to strict priority, three other scheduling methods can be configured along with a scheduler weight. The scheduler weight determines the priority of the queue in relation to the other queues in the egress port queue group. When configured for any of these scheduler modes, if the scheduling weight of a queue in the egress port queue group is set to 0, the scheduler will treat that queue as if it is in strict priority mode.

- Weighted Round Robin (WRR) - queues in the queuing group are scheduled in a weighted fashion according to the per-queue scheduler-weight.
- Round-Robin (RR) - queues in the queuing group are scheduled as if all queues in the queuing group have equal weighting.
- Weighted Deficit Round Robin (WDRR) - queues in the queuing group are scheduled in a weighted fashion according to the per-queue scheduler-weight and frame size. When using this scheduler algorithm, you can configure the scheduler granularity (wdrr-scheduler-granularity) to adjust pool sizes in which the weights are applied. This setting is especially useful when using port speeds lower than 10Mbps. Small pool sizes have less latency, and more jitter. Larger sizes have more latency, and less jitter. The value entered will be rounded to the nearest largest size of 50, 100, 200, 400, 800, or 1600 Kbytes. The default value is 1600 Kbytes.

  Note: When egress scheduling is set to either WRR or WDRR on 3916, 3930, 3931, or 3960 devices, the CIR and CBS settings are ignored. Do not configure CIR and CBS for WRR and WDRR scheduling modes. Instead, configure the weight values to distribute the traffic associated with each queue.

The default scheduler weights per queue are as follows:

- Queue 0: 20
- Queue 1: 30
- Queue 2: 40

Copyright© 2012-2013 Ciena® Corporation

- Queue 3: 50
- Queue 4: 60
- Queue 5: 70
- Queue 6: 80
- Queue 7: 0

Egress shaping

A shaper-rate and burst size can be configured on the egress port queue group that limits the amount of egress bandwidth (in Kbps) a port uses(0 to 40 Gbps) and controls the amount of traffic (in Kbytes) that can burst (0 to 256 Mbytes). The configured shaper-rate normalizes to a rate in varying increments depending upon the sum of the configured burst size.

In order to guarantee CIR, the individually-set CIR on queues must not be oversubscribed for the egress port that the frames are egressing. That is, the sum of all CIRs for the physical queues and the virtual queues in a queue group on a port should not exceed the administrative (and operational) speed of the port. When the shaper rate of the egress port queue group is set, it sets a single Information Rate (IR) regardless of the CIR and EIR of the individual queues. To guarantee CIR, the sum of the CIRs on each queue cannot exceed the IR of the egress port queue group.

When a queue or an attribute of the egress queue group of an aggregated port is modified, that value is updated at each queue or physical port of the aggregation, with the exception of the shaper rate bandwidth, which differs depending upon the platform:

- 3911, 3920, 3940, 5140 - The shaper rate bandwidth is divided equally among the number of distributing ports. This means that the shaper rate can be used to control the total bandwidth that the ports share, otherwise the ports would each have an amount of bandwidth that would add up to a cumulative amount of bandwidth for the aggregation member ports. The default shaper rate is the same as the administrative port speed.

  Note: 3902 switches do not support link aggregation.

- 3916, 3930, 3931, 3960, 5150 - The shaper rate bandwidth for the link aggregation defaults to a bandwidth equal to the maximum number of 10 Gig ports allowed in a link aggregation, regardless if the link aggregation consists of 1 Gig or 10 Gig ports.

For example, on a 3960 device, the maximum number of 10 Gig ports allowed in a link aggregation is 2; therefore, the default Egress Shaper-Rate Bandwidth for the link aggregation is 20 Gbps.

Copyright© 2012-2013 Ciena® Corporation

For a 5150 device, the maximum number of 10 Gig ports allowed in a link aggregation is 4; therefore, the default Egress Shaper-Rate Bandwidth for the link aggregation is 40 Gbps.

When you set the burst size, the value is rounded to the nearest larger 4Kbyte block. The burst size is not divided among ports in an aggregation group. The default burst size is 10240 Kbytes.

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-1
Configuring Class of Services (CoS) policies on a port

You can configure CoS policies per port.

Note: The egress frame CoS policy setting is not available on 3902, 3911, 3920, 3940, and 5140 devices.

The resolved CoS Policy options are:

- Outer .1D mapped (dot1d-tag1-cos) -The Priority Code Point (PCP)/Layer 2 (L2) CoS 802.1D priority and Discard Eligibility Indicator (DEI)/ Canonical Format Indicator (CFI) value from the outer-tag is mapped to an R-CoS and R-Color values derived from the Resolved CoS Mapping table. When this policy is applied, CoS mapping is in trusted mode. This is the default for all ports.
- Fixed CoS (fixed-cos), Fixed R-CoS (fixed-rcos) and Fixed R-COLOR (fixed-rcolor) - CoS values in the frame are ignored and fixed R-CoS is applied to the frame from the Fixed Resolved CoS and Fixed Resolved Color value for the port. When this policy is applied, CoS mapping is in untrusted mode. The default Fixed Resolved CoS value is 0 and the default Fixed Resolved Color is green.
- Layer 3 DSCP CoS (l3-dscp-cos) - The Differentiated Services Code Point (DSCP) CoS value from the Layer 3 Type Of Service (TOS) byte is mapped to an R-CoS and R-Color value from the Resolved CoS Mapping table. When this policy is applied, CoS mapping is in trusted mode.

The egress frame CoS policy options are to ignore or to re-mark the L2 frame's PCP and DEI fields based on the internal R-CoS and R-Color values (rcos-to-l2-outer-pcp-map).

Step	Action

To set the CoS policies and mapping

1 Set the resolved CoS policy:

```
port set port <PortNameList> [resolved-cos-policy <dot1d-
tag1-cos|fixed-cos|l3-dscp-cos>] [fixed-rcos <NUMBER: 0-
7] [fixed-rcolor <green|yellow>] [egress-frame-cos-
policy <ignore | rcos-to-l2-outer-pcp-map>]
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Example

The follwing example sets a resolved CoS policy.

```
port set port 1 resolved-cos-policy dot1d-tag1-cos
```

The following example sets the policy to a fixed resolved CoS and fixed resolved color.

```
port set port 1 resolved-cos-policy fixed-cos fixed-rcos 0 fixed-rcolor yellow
```

The following example sets an egress-frame-cos-policy.

```
port set port 1 egress-frame-cos-policy rcos-to-l2-outer-pcp-map
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-2
Displaying the CoS policy and mapping on a port

The Resolved CoS Policy, Egress Frame CoS Policy, Fixed Resolved Color, and Fixed Resolved CoS settings can be displayed for a specific port.

Step	Action
1	Display the settings: `port show port <PortName>` where <PortName> is the port that you want to display settings for.

—end—

Example

The following example shows sample output for the port show port command.

```
port show port 1

+------------------------------- PORT 1 INFO -----------------------------+
| Field                     | Admin                   | Oper                  |
+---------------------------+-------------------------+-----------------------+
| Type                      | GigEthernet                                     |
...
| Fixed Resolved CoS        | 0                       | 0                     |
| Fixed Resolved Color      | green                   | green                 |
...
| Egress Frame Cos Policy   | ignore                  | ignore                |
...
| Resolved CoS Policy       | dot1d-tag1-cos          | dot1d-tag1-cos        |
...
+---------------------------+-------------------------+-----------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-3
Creating a resolved CoS map

Table 13-2 shows the default resolved CoS map (DefaultFcosRcos).

Table 13-2
Default resolved CoS map

Frame PCP/L2 CoS 802.1D	DEI/CFI	Frame L3 DSCP CoS	R-CoS	R-COLOR
0	0	0-7	0	Green
0	1	0-7	0	Green
1	0	8-15	1	Green
1	1	8-15	1	Green
2	0	16-23	2	Green
2	1	16-23	2	Green
3	0	24-31	3	Green
3	1	24-31	3	Green
4	0	32-39	4	Green
4	1	32-39	4	Green
5	0	40-47	5	Green
5	1	40-47	5	Green
6	0	48-55	6	Green
6	1	48-55	6	Green
7	0	56-63	7	Green
7	1	56-63	7	Green
Untagged	Untagged	Untagged	Ingress port fixed PCP/L2 CoS 802.1D value	Ingress port fixed R-COLOR

In addition to the default resolved CoS map, you can configure customized maps and assign them to ports.

Copyright© 2012-2013 Ciena® Corporation

You can configure 1 custom resolved CoS map per port, including the following types of mapping:

- L2 CoS - maps the customer PCP/L2 CoS 802.1D and DEI/CFI priority to R-CoS and R-COLOR.
- L3 CoS - maps the customer L3 CoS DSCP to R-CoS and R-COLOR.

Step	Action
1	Create a custom resolved CoS map: `traffic-services cos-mapping resolved-cos-map create {cos-map <Resolved CoS Map Name>}` where cos-map <Resolved CoS Map Name> is the identifier for the custom resolved CoS map.

—end—

Example

The following example creates a custom resolved CoS map with an identifier of custom-1.

```
traffic-services cos-mapping resolved-cos-map create cos-map custom-1
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-4
Modifying a resolved CoS map

You can modify the attributes for the default and all custom resolved CoS maps.

Step	Action
1	Modify a custom resolved CoS map: `traffic-services cos-mapping resolved-cos-map set {cos-map <String[15]>} [dot1dpri-cos <0-7>] [dot1dpri-dei <NUMBER:0-1>] [l3-cos-dscp <NUMBER:0-63>] [mpls-tc <String>] {r-cos <NUMBER:0-7>} [r-color <green|yellow>]` where cos-map <String[15]> is the identifier for the custom resolved CoS map. [mpls-tc <String>] overloads the equivalent dot1dpri value with dot1dei=0 entry. The mpls-tc values apply only to MPLS enabled devices. **—end—**

Example

The following example modifies the attributes of the resolved CoS map custom-1.

```
traffic-services cos-mapping resolved-cos-map set cos-map custom-1 dot1dpri-cos 1,2,3,4 r-cos 1
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-5
Setting the resolved CoS map for a port

By default, ports are set to use the default Resolved CoS Map. You can modify the port to use a custom resolved CoS map or to reset to the default. In addition, you can configure the port to remark the frame's SVLAN Layer 2 CoS value with the mapped R-CoS and R-COLOR values mapped from the Customer DSCP or 802.1D priority.

Note: Remarking the frame's Layer 3 DSCP CoS value with a mapped DSCP value is implemented with traffic profile configuration.

Step	Action
1	Set the resolved CoS map for a port: `port set port <PortNameList> resolved-cos-map <Resolved CoS Map Name> resolved-cos-remark-l2 <false|true>` where <PortNameList> is the port that you want to set a resolved CoS map for. <Resolved CoS Map Name> is the identifier for the resolved CoS map. **—end—**

Example

The following example sets the resolved CoS map for a port.

```
port set port 1 resolved-cos-map custom-1 resolved-cos-remark-l2 true
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-6
Deleting a custom R-CoS map

You can delete custom resolved CoS maps, but not the default.

Note: A custom R-CoS queue map cannot be deleted if it is assigned to a port.

Step	Action

1	Set the port using the custom R-CoS map to the default R-CoS map:
	`port set port <PortName> ingress-to-egress-qmap Default-RCOS`
	`traffic-services queuing queue-map delete rcos-map custom-map-1`

2	Delete a resolved CoS map: `traffic-services cos-mapping resolved-cos-map delete cos-map <Resolved CoS Map Name>` where <Resolved CoS Map Name> is the resolved CoS map to delete.

—end—

Example

The following example deletes the custom resolved CoS map name custom-1.

```
port set port 9 ingress-to-egress-qmap Default-RCOS
traffic-services queuing queue-map delete rcos-map custom-map-1
traffic-services cos-mapping resolved-cos-map delete cos-map custom-1
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-7
Displaying resolved CoS maps

You can

- display a summary of resolved CoS maps
- display details for resolved CoS maps

Step	Action

To display a summary of resolved CoS maps

1 Display a summary of resolved CoS maps:

```
traffic-services cos-mapping resolved-cos-map show [cos-map <cos-map>]
```

where

cos-map <cos-map> is the FCOS->RCOS name.

To display details for resolved CoS maps

2 Display details for resolved CoS maps:

```
traffic-services cos-mapping resolved-cos-map show [cos-map <cos-map>] detail
```

where

cos-map <cos-map> is the FCOS->RCOS name.

—end—

Example

The following example shows sample output for displaying resolved CoS maps.

```
traffic-services cos-mapping resolved-cos-map show cos-map custom-1
```

```
+-------------------------------------------------------------------------+
|                  Ingress COS->RCOS Map Summary - custom-1               |
+-------------------------------------------------------------------------+
|                            Ingress COS                          || RCOS | RCOLOR |
| L2 COS        | DEI/CFI ||          L3 DSCP                     ||      |        |
+---------------+---------++--------------------------------------++------+--------+
|               |         ||                                      ||      |        |
| 0             |    0    ||                                      ||    0 | green  |
| 0             |    1    ||                                      ||      | green  |
|               |         || 0-7                                   ||      | green  |
+---------------+---------++--------------------------------------++------+--------+
|               |         ||                                      ||      |        |
| 1-4           |    0    ||                                      ||    1 | green  |
```


Copyright© 2012-2013 Ciena® Corporation

```
| 1             |    1    ||                                 ||       | green  |
|               |         || 8-15                            ||       | green  |
+---------------+---------++---------------------------------++------+--------+
|               |         ||                                 ||      |        |
| 2             |    1    ||                                 ||    2 | green  |
|               |         || 16-23                           ||      | green  |
...
+---------------+---------++---------------------------------++------+--------+
|               |         ||                                 ||      |        |
| 7             |    0    ||                                 ||    7 | green  |
| 7             |    1    ||                                 ||      | green  |
|               |         || 56-63                           ||      | green  |
+---------------+---------++---------------------------------++------+--------+
```

The following example shows sample output for details of a resolved CoS map.

```
traffic-services cos-mapping resolved-cos-map show cos-map custom-1 detail
```

```
+--------------------------------------+
|     RESOLVED COS MAP DETAIL INFO     |
+--------------------------------------+
| Name                | custom-1       |
| logical-id          | 2              |
+---------------------++----------------+
| L2 Cos  | DEI/CFI |||  RCOS | RCOLOR |
+---------------------++-------+--------+
|    0    |    0    |||    0  |  green |
|    0    |    1    |||    0  |  green |
|    1    |    0    |||    1  |  green |
|    1    |    1    |||    1  |  green |
...
|    6    |    0    |||    6  |  green |
|    6    |    1    |||    6  |  green |
|    7    |    0    |||    7  |  green |
|    7    |    1    |||    7  |  green |
+---------------------++-------+--------+
|      L3 DSCP        |||  RCOS | RCOLOR |
+---------------------++-------+--------+
|          0          |||    0  |  green |
|          1          |||    0  |  green |
|          2          |||    0  |  green |
|          3          |||    0  |  green |
...
|         63          |||    7  |  green |
+---------------------++-------+--------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-8
Creating a frame CoS map

With frame CoS mapping, the R-CoS and R-Color associated with a frame is mapped to specific CoS values in the frame based upon the default Frame CoS mapping table (DefaultRcosFcos) as shown in Table 13-3.

Table 13-3
Default Frame CoS mapping table

R-CoS	R-COLOR	Frame L2 PCP (0-7)	Frame L2 DEI
0	Green	0	0
0	Yellow	0	0
1	Green	1	0
1	Yellow	1	0
2	Green	2	0
2	Yellow	2	0
3	Green	3	0
3	Yellow	3	0
4	Green	4	0
4	Yellow	4	0
5	Green	5	0
5	Yellow	5	0
6	Green	6	0
6	Yellow	6	0
7	Green	7	0
7	Yellow	7	0

In addition to the default frame CoS map, you can configure customized maps and assign them to ports.

You can configure up to 3 custom frame CoS maps to use the frame's R-CoS and R-Color to map the customer PCP/L2 CoS 802.1D and Discard Eligibility Indicator (DEI) priority values upon egress.

Copyright© 2012-2013 Ciena® Corporation

Note: Remarking the frame's Layer 3 DSCP CoS value with a mapped DSCP value is implemented with traffic profile configuration.

Step	Action
1	Create a custom frame CoS map: `traffic-services cos-mapping frame-cos-map create {cos-map <String[15]>}` where cos-map <String[15]> is the identifier for the custom frame CoS map.

—end—

Example

The following example creates a custom frame CoS map.

```
traffic-services cos-mapping frame-cos-map create cos-map custom-1
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-9
Modifying a frame CoS map

You can modify the attributes for the default and all custom frame CoS maps.

Step	Action
1	Modify a custom frame CoS map: `traffic-services cos-mapping frame-cos-map set cos-map <cos-map> {r-cos <NUMBER:0-7>} {r-color <green|yellow>} [dot1dpri-cos <NUMBER: 0-7>] {dot1dpri-dei <NUMBER:0-1>} {mpls-tc <NUMBER: 0-7>}` where cos-map <cos-map> — is the FCOS->RCOS name. r-cos <NUMBER:0-7> — is the RCOS value to use. r-color <green|yellow> — is the resolved color. dot1dpri-cos <NUMBER: 0-7> — is the egress frame COS (pcp) value. dot1dpri-dei <NUMBER:0-1> — is the egress frame color. mpls-tc <NUMBER: 0-7> — overloads the equivalent dot1dpri value with dot1dei=0 entry. The mpls-tc values apply only to MPLS enabled devices. **—end—**

Example

The following example modifies a custom frame CoS map.

```
traffic-services cos-mapping frame-cos-map set cos-map fcustom-1 r-cos 1
r-color yellow dot1dpri-cos 1 dot1dpri-dei 1
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-10
Deleting a frame CoS map

You can delete custom frame CoS maps, except for the default.

Step	Action
1	Delete a frame CoS map: `traffic-services cos-mapping frame-cos-map delete cos-map <RcosFcosMap>` where cos-map <cos-map> is the FCOS->RCOS name.

—end—

Example

The following example deletes a frame CoS map.

```
traffic-services cos-mapping frame-cos-map delete cos-map custom-1
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-11
Setting the frame CoS map for a port

By default, ports are set to ignore frame CoS mapping. You can modify the port to use a custom frame CoS map or the default. In addition, you can configure the port to remark the frame's SVLAN Layer 2 CoS and R-Color value with the mapped R-CoS value mapped from the Customer DSCP or 802.1D priority.

Note: Remarking the frame's Layer 3 DSCP CoS value with a mapped DSCP value is implemented with traffic profile configuration. Note: When egress Frame CoS mapping is enabled, the ingress setting resolved-cos-remark-l2 is overridden and will no longer preserve the incoming frame CoS.

Step	Action
1	Set the Frame CoS map for a port: `port set port <PortNameList> frame-cos-map <Frame CoS Map Name> egress-frame-cos-policy <ignore|rcos-to-l2-outer-pcp-map> resolved-cos-remark-l2 <false|true>`

—end—

Example

The following example sets the frame CoS map for a port.

```
port set port 1 frame-cos-map custom-1 egress-frame-cos-policy rcos-to-l2-
outer-pcp-map resolved-cos-remark-l2 true
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-12
Displaying frame CoS maps

Display frame CoS maps.

Step	Action
1	Display all frame CoS maps or a specific map: `traffic-services cos-mapping frame-cos-map show [frame-cos-map <Frame CoS Map Name>]` where frame-cos-map <Frame CoS Map Name> is the specific frame CoS map to be displayed.

—end—

Example

The following example shows sample output for a frame CoS map.

```
traffic-services cos-mapping frame-cos-map show

+---------------------------------------+
| RESOLVED COS MAP TO FRAME COS INFO    |
+---------------------------------------+
| Name                |  DefaultRcosFcos  |
| logical-id          |  1                |
+-------+---------++---------+---------+
|  RCOS | RCOLOR  |||  L2-COS |  L2-DEI |
+-------+---------++---------+---------+
|     0 |  green  |||     0   |     0   |
|     0 |  yellow |||     0   |     0   |
|     1 |  green  |||     1   |     0   |
|     1 |  yellow |||     1   |     0   |
|     2 |  green  |||     2   |     0   |
|     2 |  yellow |||     2   |     0   |
|     3 |  green  |||     3   |     0   |
|     3 |  yellow |||     3   |     0   |
|     4 |  green  |||     4   |     0   |
|     4 |  yellow |||     4   |     0   |
|     5 |  green  |||     5   |     0   |
|     5 |  yellow |||     5   |     0   |
|     6 |  green  |||     6   |     0   |
|     6 |  yellow |||     6   |     0   |
|     7 |  green  |||     7   |     0   |
|     7 |  yellow |||     7   |     0   |
+-------+---------++---------+---------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-13
Configuring ingress R-CoS to egress queue mapping

Each physical port supports 8 CoS queues. Depending upon the policies applied to the port, the R-CoS and R-COLOR values may change during processing. The final internal R-CoS value is used to map the frame to one of the 8 physical port CoS queues as shown in Table 13-4

Table 13-4
Default internal R-CoS mapping

R-CoS	Physical Port CoS Queue
0	0
1	0
2	1
3	2
4	3
5	4
6	5
7	6
Reserved for outbound CPU frames	7

In addition to the default queue mapping, you can create up to 7 custom Queue Map Profiles map R-CoS to CoS queue. The mapping-mode of custom queue map profiles is always R-CoS mapped. To configure a port to use a custom queue map profile, set the *ingress-to-egress-qmap* attribute. R-CoS values (0 - 7) can be mapped to CoS Queues (0 - 7), or any combination thereof (for example, all 7 R-CoS values mapped to the same CoS queue).

To create and modify a custom R-CoS to CoS queue table and assign it to an egress port:

```
traffic-services queuing queue-map create rcos-map custom-map-1
traffic-services queuing queue-map set rcos-map custom-map-1 rcos 0 queue 3
port set port 9 ingress-to-egress-qmap custom-map-1
```

To set the R-CoS map to the default for a port:

Note: You cannot unset a port's ingress-to-egress-qmap profile, as this is a required parameter. To remove a custom R-CoS profile from a port you must replace it with another R-CoS profile (for example, the Default-RCOS profile).

```
port set port 9 ingress-to-egress-qmap Default-RCOS
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-14
Displaying an R-CoS map

Display an R-Cos map.

Step	Action

1 To display an R-CoS map:

```
traffic-services queuing queue-map show [rcos-map
<RcosQueueMap>]
```

where

rcos-map <RcosQueueMap>	is the R-CoS map to be displayed.

—end—

Example

The following example shows output for the default R-CoS map.

```
traffic-services queuing queue-map show Default-RCOS

+---------------------------RCOS TO QUEUE MAPPING ------------------------+
|                                                                          |
+-------------------------------+------------------------------------------+
| Name                          | Default-RCOS                             |
| Id                            | 1                                        |
| Type                          | RCOS To Queue Mapping                    |
+-------------------------------------------------------------------------+
| RCOS:    | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |                                  |
|          +---+---+---+---+---+---+---+---+                              |
| Queue:   | 0 | 0 | 1 | 2 | 3 | 4 | 5 | 6 |                                  |
+---------+---+---+---+---+---+---+---+---+---------------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-15
Applying R-CoS policies and mapping in a VLAN

With the example in this section, frames with VID 100 and specific PCP/L2 CoS 802.1D and DEI/CFI values are assigned an internal R-CoS as shown in Table 13-5.

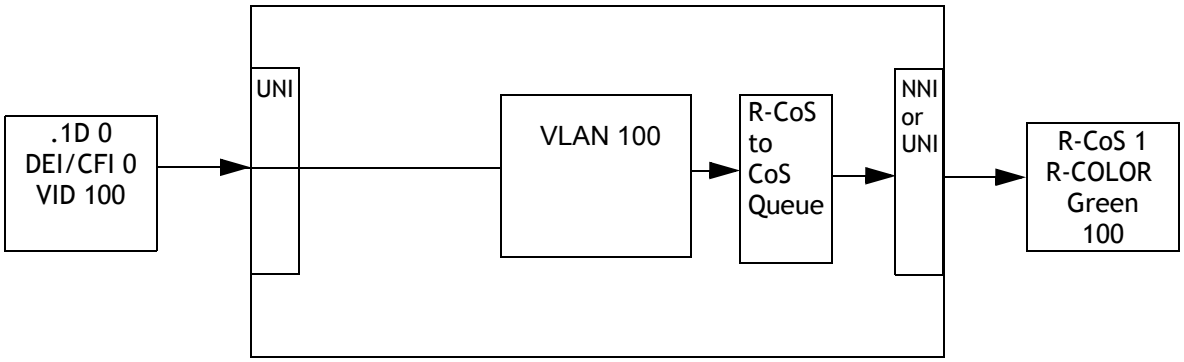
Table 13-5
Custom map

Frame PCP/ L2 CoS 802.1D	Frame DEI/ CFI	R-CoS	R-COLOR
0	0,1	1	Green
1	0,1	0	Yellow
2	0,1	1	Green
3	0,1	1	Green
4	0,1	2	Green
5	0,1	3	Green
6	0,1	4	Green
7	0,1	0	Yellow

Untagged frames will be assigned an R-CoS according to the port's fixed R-CoS configuration. For example, a frame with VID 100 and PCP/L2 CoS.1D value of 0 is mapped to R-CoS 1.

Copyright© 2012-2013 Ciena® Corporation

Figure 13-3



Step	Action

1 Create the R-CoS map and set the mapping for PCP/L2 CoS.1D to R-CoS values.

```
traffic-services cos-mapping resolved-cos-map create cos-map MyL2Map
traffic-services cos-mapping resolved-cos-map set cos-map MyL2Map dot1dpri-
cos 0,2,3 dot1dpri-dei 0,1 r-cos 1 r-color green
traffic-services cos-mapping resolved-cos-map set cos-map MyL2Map dot1dpri-
cos 1,7 dot1dpri-dei 0,1 r-cos 0 r-color yellow
traffic-services cos-mapping resolved-cos-map set cos-map MyL2Map dot1dpri-
cos 4 dot1dpri-dei 0,1 r-cos 2 r-color green
traffic-services cos-mapping resolved-cos-map set cos-map MyL2Map dot1dpri-
cos 5 dot1dpri-dei 0,1 r-cos 3 r-color green
traffic-services cos-mapping resolved-cos-map set cos-map MyL2Map dot1dpri-
cos 6 dot1dpri-dei 0,1 r-cos 4 r-color green
```

2 Configure the ingress port's resolved CoS Policy and Map.

```
port set port 2 resolved-cos-policy dot1d-tag1-cos resolved-cos-map MyL2Map
```

3 Create the provider VLANs and add the NNI port.

```
vlan create vlan 100,1000-1001
vlan add vlan 1000-1001 port 6
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-16
Applying R-CoS policies and mapping in a virtual switch

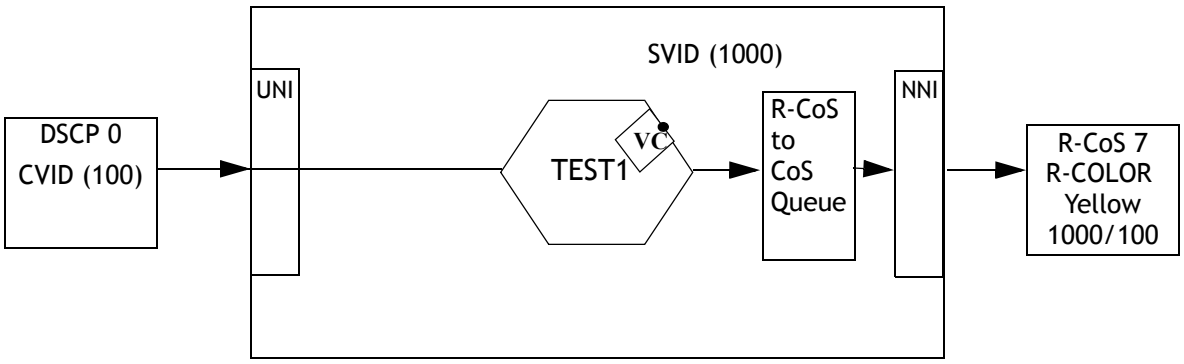
With the example in this section, frames with CVID 100 and specific DSCP values are assigned an internal R-CoS as shown in Table 13-6, because the virtual switch member inherits the custom R-CoS map of the port. Untagged frames will be assigned an R-CoS according to the port's fixed R-CoS configuration.

Table 13-6
Custom map

Frame DSCP	R-CoS	R-COLOR
6	1	Green
5	2	Green
4	3	Green
3	4	Green
2	5	Green
1	6	Green
0, 8	7	Yellow

As an example, a frame with CVID 100 and DSCP value of 0 is mapped to R-CoS 7 and R-COLOR Yellow.

Figure 13-4



Copyright© 2012-2013 Ciena® Corporation

Step	Action

1 Create the R-CoS map and set the mapping for DSCP to R-CoS values.

```
traffic-services cos-mapping resolved-cos-map create cos-map MyDscpMap
traffic-services cos-mapping resolved-cos-map set cos-map MyDscpMap l3-cos-
dscp 6 r-cos 1 r-color green
traffic-services cos-mapping resolved-cos-map set cos-map MyDscpMap l3-cos-
dscp 5 r-cos 2 r-color green
traffic-services cos-mapping resolved-cos-map set cos-map MyDscpMap l3-cos-
dscp 4 r-cos 3 r-color green
traffic-services cos-mapping resolved-cos-map set cos-map MyDscpMap l3-cos-
dscp 3 r-cos 4 r-color green
traffic-services cos-mapping resolved-cos-map set cos-map MyDscpMap l3-cos-
dscp 2 r-cos 5 r-color green
traffic-services cos-mapping resolved-cos-map set cos-map MyDscpMap l3-cos-
dscp 1 r-cos 6 r-color green
traffic-services cos-mapping resolved-cos-map set cos-map MyDscpMap l3-cos-
dscp 0,8 r-cos 7 r-color yelllow
```

2 Configure the ingress port's resolved CoS Policy and Map.

```
port set port 2 resolved-cos-policy l3-dscp-cos resolved-cos-map MyDscpMap
```

3 Create the provider VLANs and add the NNI port.

```
vlan create vlan 100,1000
vlan add vlan 1000 port 6
```

4 Create the virtual circuits.

```
virtual-circuit ethernet create vc TEST1 vlan 1000 statistics on
```

5 Reserve VLANs for the virtual switches.

```
virtual-switch add reserved-vlan 4090-4091
virtual-switch ethernet create vs TEST1 vc TEST1
```

6 Add the members with the encap policy to inherit from the port settings.

```
virtual-switch ethernet add vs TEST1 port 2 encap-cos-policy port-inherit
```

7 Set the untagged data virtual switch for the member port.

```
port set port 2 untagged-data-vs TEST1
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-17
Enabling traffic profiling

By default, traffic profiling status is disabled globally and per port. In order to configure traffic profiling, EIR/PIR, and CIR values must be provisioned.

Step	Action
1	Enable traffic profiling: `traffic-profiling enable [<PortNameList>]` where <PortNameList> is the port that you want to enable traffic profiling on.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-18
Setting the traffic profiling provisioning mode

Set the provisioning mode. You can set the provisioning mode to

- EIR
- PIR

Step	Action

To set the provisioning mode to EIR

1 Set the provisioning mode to EIR:

```
traffic-profiling set meter-provisioning eir
```

To set the provisioning mode to PIR

2 Set the provisioning mode to PIR:

```
traffic-profiling set meter-provisioning pir
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-19
Displaying traffic profiling information

You can display:

- Status global and port
- Per port attributes
- Classification mode for all ports
- Meter pool assignments for all ports on 3902, 3911, 3920, 3940, and 5140 platforms

Step	Action

To display traffic profiling status

1 Display traffic profiling status:

```
traffic-profiling show
```

2 Display per port attributes:

```
traffic-profiling show port <PortName>
```

where

port <PortName> is the port that you want to display attributes for.

3 Display the classification mode for all ports:

```
traffic-profiling show classifier-mode
```

4 Display meter pool assignments for all ports:

```
traffic-profiling show meter-pool
```

—end—

Example

The following example shows the traffic-profiling status.

```
traffic-profiling show
+------------ TRAFFIC PROFILING GLOBAL TABLE ---------------+
| Profiling Status      | Disabled                           |
| meter-provisioning    | eir                                |
+-----------------------------------------------------------+

+---------------- PORT TRAFFIC PROFILE TABLE ----------------+
|  Port      |        Status         | Mode                  |
|            | Admin     | Oper      |                       |
+------------+-----------+-----------+-----------------------+
| 1          | Disabled  | Disabled  | standard-dot1dpri     |
| 2          | Disabled  | Disabled  | standard-dot1dpri     |
| 3          | Disabled  | Disabled  | standard-dot1dpri     |
| 4          | Disabled  | Disabled  | standard-dot1dpri     |
| 5          | Disabled  | Disabled  | standard-dot1dpri     |
```


Copyright© 2012-2013 Ciena® Corporation

```
| 6          | Disabled | Disabled | standard-dot1dpri         |
| 7          | Disabled | Disabled | standard-dot1dpri         |
| 8          | Disabled | Disabled | standard-dot1dpri         |
| 9          | Disabled | Disabled | standard-dot1dpri         |
| 10         | Disabled | Disabled | standard-dot1dpri         |
| 11         | Disabled | Disabled | standard-dot1dpri         |
| 12         | Disabled | Disabled | standard-dot1dpri         |
+-----------+----------+----------+---------------------------+
```

The following example shows the traffic profiling per port attributes.

```
traffic-profiling show port 1
+---------------- TRAFFIC PROFILING TABLE -------------------+
| Port                          | 1                         |
| Profiling Admin State         | Disabled                  |
| Profiling Oper State          | Disabled                  |
| Profiling Mode                | standard-dot1dpri         |
| Meter Pool                    | TP-POOL3                  |
| Non-Conform Standard Profile  | drop                      |
| ARP Standard Profile          | bypass                    |
| Classification Mode           | narrow                    |
+-------------------------------------------------------------+
```

The following example displays the classification mode for all ports.

```
traffic-profiling show classifier-mode
+------ TRAFFIC-PROFILING PORT CLASSIFIER MODE -----+
| Port            |         Port Class Mode          |
+-----------------+----------------------------------+
| 1               |         wide                     |
| 2               |         narrow                   |
...
+-----------------+----------------------------------+
```

The following example shows meter pool assignments on a 3902 device.

```
traffic-profiling show meter-pool
+----------- TRAFFIC-PROFILING METER-POOL MAP ----------+
|    Meter-Pool    |        Ports         |Total | Used |
+------------------+----------------------+------+------+
| TP-POOL1         |                      | 64   | 0    |
+------------------+----------------------+------+------+
| TP-POOL2         |                      | 64   | 0    |
+------------------+----------------------+------+------+
| TP-POOL3         |  1 2                 | 64   | 0    |
+------------------+----------------------+------+------+
| TP-POOL4         |  3 4                 | 64   | 0    |
+------------------+----------------------+------+------+
| TP-POOL5         |  5 6                 | 64   | 0    |
+------------------+----------------------+------+------+
| TP-POOL6         |  7 8                 | 64   | 0    |
+------------------+----------------------+------+------+
| TP-POOL7         |  9 10                | 64   | 0    |
+------------------+----------------------+------+------+
| TP-POOL8         |  11 12               | 64   | 0    |
+------------------+----------------------+------+------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-20
Setting traffic profiling port attributes

Traffic profiling can be configured per port and per EVC with the following attributes:

- arp-standard-profile - Sets whether non-conforming ARP frames are associated with an existing standard profile or bypass traffic profiling. ARP frames are treated specially because they are required for IP networks to function. Dropping ARP frames would result in breaking an IP network.The default setting is bypass.
- meter-pool - A meter pool provides the meter resources for traffic profiles on 3902, 3911, 3920, 3940, and 5140 platforms. For 3916, 3930, 3931, 3960, and 5150 platforms, the meter-pool attribute does not apply. These platforms use the global meter pool. The number of supported meters per meter pool is 64. A single traffic profile consumes 1 meter in the meter pool, and 2 hardware meter resources. You can change the meter pool for a given port as long as there are enough resources in the newly assigned meter pool. A given port can only be assigned to one meter pool. Default metering pool assignments are shown in Table 13-7.

  Note: The number of traffic profiles that can be supported on 3902, 3911, 3920, 3940, and 5140 platforms is less than the number of meters per pool. The actual number of traffic profiles per metering pool depends upon the number of ports in the meter pool and the number of classifiers configured per profile. For additional information regarding traffic profiling resources, see “Hardware resource management” on page 4-1.

Table 13-7
Metering pools default port assignment per platform

Platform	Meter pool	Ports
3902	TP-POOL1	1
	TP-POOL2	2

Copyright© 2012-2013 Ciena® Corporation

Table 13-7
Metering pools default port assignment per platform

Platform	Meter pool	Ports
3911	TP-POOL1	None
	TP-POOL2	None
	TP-POOL3	None
	TP-POOL4	1, 2
	TP-POOL5	3, 4
	TP-POOL6	5, 6
	TP-POOL7	7, 8
	TP-POOL8	9, 10
3920	TP-POOL1	None
	TP-POOL2	None
	TP-POOL3	1, 2
	TP-POOL4	3, 4
	TP-POOL5	5, 6
	TP-POOL6	7, 8
	TP-POOL7	9, 10
	TP-POOL8	11, 12
3940, 5140	TP-POOL1	1, 2, 3
	TP-POOL2	4, 5, 6
	TP-POOL3	7, 8, 9
	TP-POOL4	10, 11, 12
	TP-POOL5	13, 14, 15
	TP-POOL6	16, 17, 18
	TP-POOL7	19, 20, 21
	TP-POOL8	22, 23, 24

- classifier-mode - The classifier mode determines whether frames that are to be dropped at ingress are applicable to the meter rates of all traffic profiles configured against the port. When the mode is set to "narrow," traffic profile classification is restricted to only match frames that will not

Copyright© 2012-2013 Ciena® Corporation

be dropped at ingress. When the mode is set to "wide", traffic profile classification is not restricted and may include traffic that is to be dropped at ingress. The default classifier-mode is narrow.

- mode - The mode sets the inspection point for classification:
  - none - Removes the default classifiers for the port freeing them to be used by other ports in the same meter pool to use those resources. Before setting the mode to “none,” remove any traffic profiles created for the port. Traffic profiles cannot be created for a port that is in the “none” mode. Applicable for 3902, 3911, 3920, 3940, and 5140 platforms. 3916, 3930, 3931, 3960, and 5150 platforms consume classifiers from the global meter pool.
  - advanced - Supports EVC based classification, and supports standard and hierarchical metering simultaneously. Not applicable for 3911, 3920, and 3940 platforms. Note that non-conforming and ARP profiles are not applicable to a port in advanced mode. Ports in advanced mode behave like ports in hierarchical-vlan mode.
  - standard-dot1dpri - Finds a matching traffic profile using 802.1D priority. This mode is the default.
  - standard-ip-prec - Finds a matching profiling using the upper 3 bits of the TOS byte that make up the IP precedence.
  - standard-dscp - Finds a matching profiling using the DSCP value.
  - standard-vlan - Finds a matching profiling based upon the VLAN ID.
  - standard-vlan-dot1dpri - Finds a matching profiling based upon the VLAN ID and outer.1D priority value in the frame.
  - standard-vlan-ip-prec - Finds a matching profiling based upon the VLAN ID and ip-precedence value in the frame.
  - standard-vlan-dscp - Finds a matching profiling based upon the VLAN ID and the DSCP value in the frame.
- nonconform-standard-profile - Sets whether non-conforming frames (that are not ARP) are associated with a standard profile or dropped. Non-conforming frames are those ingressing frames that do not match one of the configured per-port traffic profiles. For example, non-conforming 802.1D priority frames are those frames that do not have a matching 802.1D priority configured in one of the traffic profiles. The same applies to DSCP (no matching 6 bit code point), and IP-Precedence (no matching 3 bit IP Precedence value). Note that certain frames, such as ARPs and non-IP frames, will always be non-conforming because these frames do not contain the DSCP field. The default setting is drop.

Copyright© 2012-2013 Ciena® Corporation

Note: The classifier-mode attribute does not apply to the nonconform-standard-profile.

Step	Action
1	Set the per port attributes: `traffic-profiling set port <PortName> {[arp-standard-profile <bypass|<Traffic ProfilingStandardName>], [meter-pool <TpMeterPoolName>], [classifier-mode <narrow|wide>], [mode <none|advanced|standard-dot1dpri|standard-ip-prec|standard-dscp|standard-vlan|standard-vlan-dot1dpri|standard-vlan-ipp|standard-vlan-dscp|hierarchical-port|hierarchical-vlan], [nonconform-standard-profile <drop|<TrafficProfilingStandardName>>]}` **—end—**

Example

The following example sets per-port attributes.

```
traffic-profiling set port 1 arp-standard-profile bypass meter-pool TP-POOL1
mode standard-vlan nonconform-standard-profile drop
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-21
Configuring a traffic profiling standard profile

You can perform the following operations on a standard profile:

- create
- modify
- remove classification attributes
- delete

If a port's ARP and non-conforming standard profiles are set to the standard profile to be deleted, set them to the default settings.

Traffic profile remarking policies take precedence over the RCOS->FCOS map settings configured on the egress port.

Step	Action

To create a standard profile

1 Create a standard profile:

```
traffic-profiling standard-profile create {port <port>}
[profile <NUMBER: 1-4096>] [dot1dpri <NUMBER>] [ip-prec
<NUMBER>] [dscp <NUMBER>] [dscp-remark-policy <leave |
fixed>] [fixed-dscp <NUMBER: 0-63>] {[cir <NUMBER>], [eir
<NUMBER>],[pir <NUMBER>], [cbs <NUMBER: 0-262144>], [ebs
<NUMBER: 0-262144>]} [name <String[15]>] [vlan <VLAN>][vs
<Virtual Switch Ethernet Name>] [statistics <on|off>]
[untagged] [parent <Traffic Profiling Standard Name>]
[child-mode <standard-dot1dpri | standard-ip-prec |
standard-dscp | standard-vlan | vlan-cos>] [default-
profile <true|false>] drop <true | false> [ingress-color-
aware <on|off>] [remark-color-policy <leave|yellow-to-
green|green-to-yellow>] [remark-rcos-policy
<leave|remark-green|remark-yellow|remark-both] [green-
remark-rcos <NUMBER: 0-7>] [yellow-remark-rcos <NUMBER:
0-7>].
```

where

{port <port>}	is the port number.
[profile <NUMBER: 1-4096>]	is the traffic profile number.
[dot1dpri <NUMBER>]	is the 802.1D priority value of matching frames.

Copyright© 2012-2013 Ciena® Corporation

where	
[ip-prec <NUMBER>]	is the IP Precedence value of matching frames.
[dscp <NUMBER>]	is the Differentiated Services Code Point (DSCP) value of matching frames.
[dscp-remark-policy <leave | fixed>]	is the DSCP Remark Policy. • leave: leave as is • fixed: IPv4 frames that classify to the traffic profile have the DSCP value marked with a fixed value,
[fixed-dscp <NUMBER: 0-63>]	is the value to use when remarking an IPv4 frame for a traffic profile with a dscp-remark-policy of fixed.
{[cir <NUMBER>]	is the committed information rate in Kbps of this traffic meter, rounded to fit hardware.
[eir <NUMBER>]	is the excess information rate in Kbps of this traffic meter, rounded to fit hardware.
[pir <NUMBER>]	is the peak information rate in Kbps of this traffic meter, rounded to fit hardware.
[cbs <#>],	is the meter committed burst size in Kbytes, rounded to nearest size supported in hardware.
[ebs <#>]}	is the excess burst size in Kbytes, rounded to nearest size supported in hardware.
[name <String[15]>]	is the traffic profile name. If you do not specify a name, one is automatically created for the traffic profile number, that is, STD#1 through STD#32.
[vlan <VLAN>]	is the VLAN ID to monitor.
[vs <Virtual Switch Ethernet Name>]	classifies the profile based on the Ethernet VS ID. Note that traffic profiles cannot classify to an MPLS VS.
[statistics <on|off>]	Enable or disable collection of statistics for the traffic profile
[untagged]	Untagged frame classifier
[child-mode <standard-dot1dpri| standard-ip-prec| standard-dscp| standard-vlan|vlan-cos>]	is the remarking mode for child profiles. Applicable to parent profiles for the 3916, 3930, 3931, 3960, and 5150 operating in hierarchical-port or hierarchical-vlan mode.

Copyright© 2012-2013 Ciena® Corporation

where

[default-profile <true|false>]	indicates whether the profile is to be a default profile for metering traffic that does not classify to the parent or child profiles. Applicable to hierarchical profiles for the 3916, 3930, 3931, 3960, and 5150.
[drop <true|false>]	indicates whether the profile automatically drops all traffic that classifies to this profile. ***Note:*** The very first frame counted in the statistics on a profile with the drop attribute set may be counted as "forwarded" when it is actually dropped as it takes one frame to move the hardware out of the previous profile.
[ingress-color-aware <on|off>]	indicates whether the meter is color aware or color blind.
[remark-color-policy <leave|yellow-to-green|green-to-yellow>]	is the meter output color remarking policy. Policies are: • leave: leaves the R-Color of the frame as it was classified by the port. This command is useful if the remark-color-policy was changed to yellow-to-green or green-to-yellow and needs to be changed back to the default. • yellow-to-green: changes the color of frames that were classified by the port as yellow to green. • green-to-yellow: changes the color of frames that were classified by the port as green to yellow. ***Note:*** This attribute cannot be set on hierarchical parent profiles.

Copyright© 2012-2013 Ciena® Corporation

where	
[remark-rcos-policy <leave|remark-green|remark-yellow|remark-both]	is the meter output color-based rcos remarking policy. Policies are: • leave: leaves the R-CoS policy without change. This command is useful if the remark-rcos-policy was changed to remark-green or remark-yellow and it needs to be changed back to the default. • remark-green: assigns a new R-CoS value to any packets that are classified as green. The new R-CoS value is assigned using green-remark-rcos <NUMBER: 0-7>. • remark-yellow: assigns a new R-CoS value to any packets that are classified as yellow. The new R-CoS value is assigned using yellow-remark-rcos <NUMBER: 0-7>. • remark-both: assigns a new R-CoS value to any packets that are classified as green and to any packets that are classified as yellow. The new R-CoS value for green packets is assigned using green-remark-rcos <NUMBER: 0-7>. The new r-cos value for yellow packets is assigned using yellow-remark-rcos <NUMBER: 0-7>. ***Note:*** This attribute cannot be set on hierarchical parent profiles.
[green-remark-rcos <NUMBER: 0-7>]	is the RCOS value for green remarking.
[yellow-remark-rcos <NUMBER: 0-7>]	is the RCOS value for yellow remarking.



CAUTION
Traffic Disruption Risk
Use caution when enabling the drop attribute when only one profile is present on the port. When a profile matches remote management traffic, it can cause the device to become unreachable.

T***o modify a standard profile***

2 Modify attributes for a standard profile:

```
traffic-profiling standard-profile set {port <PortName>}
{profile <ProfileNumber>} {[dot1dpri <NUMBER LIST: 0-7>],
[ip-prec <NUMBER LIST: 0-7>], [dscp <NUMBER LIST: 0-63>],
[dscp-remark-policy <leave | fixed>], [fixed-dscp <0-
63>], [cir <NUMBER>], [pir <NUMBER>], [cbs <#>], [ebs
```


Copyright© 2012-2013 Ciena® Corporation

```
<#>], [name <TrafficProfilingStandardName[15]>], [vlan
<VLAN>], [vs <VsEth>], [statistics <on|off>] [untagged]
[drop <true|false>] [parent <TrafficProfStd>] [child-mode
<standard-dot1dpri|standard-ip-prec|standard-dscp|..>]
[default-profile <true|false>] [ingress-color-aware
<on|off>] [remark-color-policy <leave|yellow-to-
green|green-to-yellow>] [remark-rcos-policy
<leave|remark-green|remark-yellow|remark-both] [green-
remark-rcos <#>] [yellow-remark-rcos <#>]}
```

To remove classification attributes for a standard profile

3 Remove classification attributes for a standard profile:

```
traffic-profiling standard-profile unset {port
<PortNameList>} {profile
<TrafficProfilingStandardName[15]>} [dot1dpri <NUMBER
LIST: 0-7>] [ip-prec <NUMBER LIST: 0-7>] [dscp <NUMBER
LIST: 0-63>] [vlan] [untagged]
```

To delete a standard profile

4 Delete a standard profile:

```
traffic-profiling standard-profile delete {port
<PortNameList>} {profile <TrafficProfilingStandardName>}
```

—end—

Example

The following example uses EIR mode.

```
traffic-profiling set meter-provisioning eir
traffic-profiling standard-profile create port 1 name TP1 cir 256 eir 512
dot1dpri 1
```

The following example uses PIR mode.

```
traffic-profiling set meter-provisioning pir
traffic-profiling standard-profile create port 1 name TP1 cir 256 pir 768
dot1dpri 1
```

The following example sets the child mode as part of the parent profile.

```
traffic-profile standard create port AGGiTest name P1 cir 50048 child-mode
vlan-cos vs VS1
```

The following example remarks the metered output of a traffic profile.

```
traffic-profiling standard-profile create port 1 profile 1 remark-color-
policy yellow-to-green
```

The following example remarks the R-CoS per traffic-profile.

```
traffic-profiling standard-profile create port 2 profile 2 remark-rcos-policy
remark-green
```


Copyright© 2012-2013 Ciena® Corporation

The following example specifies that a traffic profile is color blind.

```
traffic-profiling standard-profile create port 1 profile 1 ingress-color-
aware off
```

The following example sets the .1D priority of matching frames to 1 for traffic profile 1 on port 1.

```
traffic-profiling standard-profile set port 1 profile 1 dot1dpri 1
```

The following example unsets the standard traffic profile attributes.

```
traffic-profiling standard-profile unset port 1 profile 1 dot1dpri 1
```

In this example, assume a stream of frames enter the ingress port with a P-bit value of 3. If no remark-rcos-policy is used, all the frames will have R-CoS value of 3, so all the frames will be stored in the same egress queue and at egress port they all will leave with P-bit of 3. To direct the Green frames and Yellow frames to different egress queues you can use the remark-rcos-policy to assign different R-CoS values to Yellow frames. But if the default R-CoS to F-CoS mapping is used, the Yellow frames in the egress port will be assigned the new P-bit values and will leave the port with the new P-bit value.

The following example assigns all Yellow frames in a port r-cos a value of 1:

```
traffic-profiling standard-profile set port <port number> profile <profile
name> remark-rcos-policy remark-yellow yellow-remark-rcos 1
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-22
Configuring per-port standard traffic profiling

The configuration example in this section shows how to configure per-port traffic profiling to meter ingress traffic based upon 802.1D priority.

Step	Action
1	Enable traffic profiling globally. `traffic-profiling enable`
2	Enable traffic profiling on port 1. `traffic-profiling enable port 1`
3	Confirm the global configuration (optional). `traffic-profiling show`
4	Set traffic profiling mode on port 1 to standard-dot1dpri (default) and a narrow classifier-mode. `traffic-profiling set port 1 mode standard-dot1dpri classifier-mode narrow`
5	Confirm the per port configuration for port 1. `traffic-profiling show port 1`
6	Create Traffic profile 1 on port 1 with CIR 3,000Kbps and PIR 3,200Kbps, dot1dpri 0,1,2. `traffic-profiling standard-profile create port 1 profile 1 cir 3000 pir 3200 dot1dpri 0,1,2`
7	Create Traffic profile 2 on port 1 with CIR 3,000Kbps and PIR 3,200Kbps, dot1dpri 3,4,5. `traffic-profiling standard-profile create port 1 profile 2 cir 3000 pir 3200 dot1dpri 3,4,5`
8	Create Traffic profile 3 on port 1 with CIR 2,000Kbps and PIR 2,432Kbps, dot1dpri 6,7. `traffic-profiling standard-profile create port 1 profile 3 cir 2000 pir 2432 dot1dpri 6,7`
9	Create Traffic profile 4 on port 1 with CIR 128Kbps and PIR 256Kbps for untagged traffic. `traffic-profiling standard-profile create port 1 profile 4 untagged cir 128 pir 256 untagged`
10	Display the traffic profiles associated with port 1 (Optional). `traffic-profiling standard-profile show port 1`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-23
Configuring per-port and per-VLAN standard traffic profiling

The configuration example in this section shows how to configure per-port and per-VLAN traffic profiling to meter ingress traffic based upon VLAN tag. Port 1 is the UNI port and port 9 is the NNI port.

Step	Action
1	Enable traffic profiling globally. `traffic-profiling enable`
2	Enable traffic profiling on UNI port 1. `traffic-profiling enable port 1`
3	Confirm the global configuration (optional). `traffic-profiling show`
4	Set traffic profiling mode on port 1 to standard-vlan and a narrow classifier-mode. `traffic-profiling set port 1 mode standard-vlan classifier-mode narrow`
5	Confirm the per port configuration for port 1. `traffic-profiling show port 1`

6	Create VLAN 1001, 1002, 1003.
	`vlan create vlan 1001-1003`
	
	`...This may take a while depending upon the VLAN range...`

7	Add port 1 and 9 to VLAN1001, 1002, 1003.
	`vlan add vlan 1001-1003 port 1,9`
	
	`...This may take a while depending upon the VLAN range...`

8	Confirm the VLAN configuration (optional). `vlan show`
9	Create Traffic profile 1 on port 1 with CIR 100 kbps and PIR 1024 Kbps, and classifier VLAN 1001. `traffic-profiling standard-profile create port 1 profile 1 cir 100 pir 1024 vlan 1001`
10	Create Traffic profile 2 on port 1 with CIR 100 kbps and PIR 1024 Kbps, and classifier VLAN 1002. `traffic-profiling standard-profile create port 1 profile 2 cir 100 pir 1024 vlan 1002`

Copyright© 2012-2013 Ciena® Corporation

11 Create Traffic profile 3 on port 1 with CIR 100 kbps and PIR 1024 Kbps, and classifier VLAN 1003.

```
traffic-profiling standard-profile create port 1 profile
3 cir 100 pir 1024 vlan 1003
```

12 Create Traffic profile 4 on port 1 with CIR 128Kbps and PIR 256Kbps for untagged traffic.

```
traffic-profiling standard-profile create port 1 profile
4 untagged cir 128 pir 256 untagged
```

13 Display the traffic profiles associated with port 1 (Optional).

```
traffic-profiling standard-profile show port 1
```

—end—

Copyright© 2012-2013 Ciena® Corporation

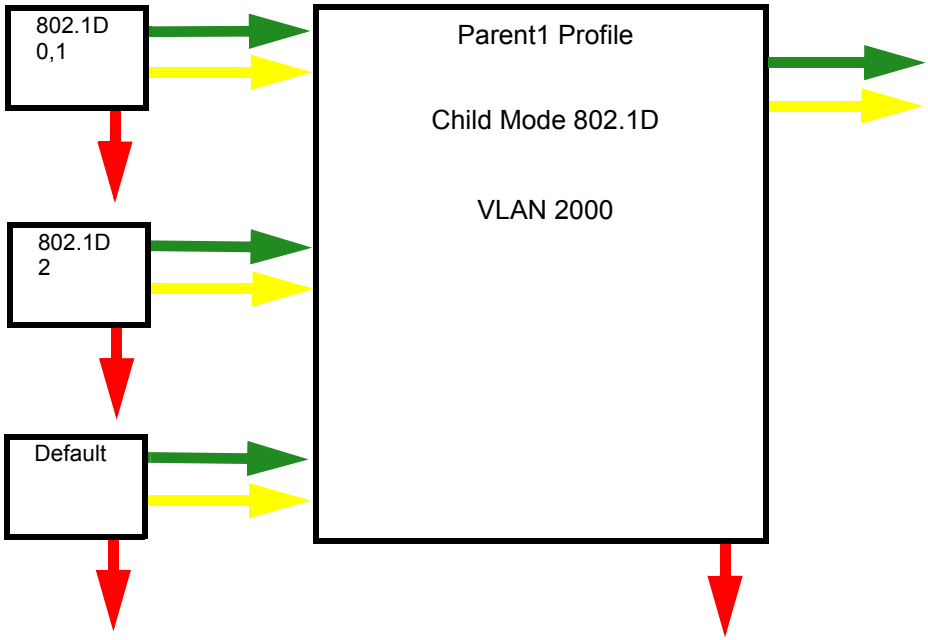
Procedure 13-24
Configuring hierarchical VLAN traffic profiles

In the configuration example for this section, three parent profiles are configured for three VLANs (2000, 2100, 4000) with different child modes and profiles for 802.1D, DSCP, and IPP classification.

Step	Action
1	Set the port mode to hierarchical VLAN mode. `traffic-profiling set port 6 mode hierarchical-vlan`
2	Create the parent profile to classify traffic with VLAN 2000 with a narrow classifier-mode and specify the child mode to classify based on 802.1D priority. `traffic-profiling standard-profile create port 6 name parent1 child-mode standard-dot1dpri cir 50000 cbs 8 vlan 2000 classifier-mode narrow`
3	Create the child profile to classify traffic with 802.1D priority values 0 and 1. `traffic-profiling standard-profile create port 6 name child1 parent parent1 dot1dpri 0,1 cir 30000 pir 50000 cbs 4 ebs 2`
4	Create the child profile to classify traffic with 802.1D priority values 2. `traffic-profiling standard-profile create port 6 name child2 parent parent1 dot1dpri 2 cir 10000 pir 50000 cbs 4 ebs 2`
5	Create a child default under vlan 2000. `traffic-profiling standard-profile create port 6 default-profile true cir 64 pir 45000 cbs 4 ebs 4 name vlan2000child parent parent1`

Copyright© 2012-2013 Ciena® Corporation

Figure 13-5
Hierarchical VLAN with 802.1D priority



6 Create the parent profile to classify traffic with VLAN 4000 with a narrow classifier-mode and specify the child mode to classify based on DSCP priority.

```
traffic-profiling standard-profile create port 6 name
parent2 child-mode standard-dscp cir 50000 cbs 8 vlan 4000
```

7 Create the child profile to classify traffic with DSCP value 8.

```
traffic-profiling standard-profile create port 6 name
child1dscp parent parent2 dscp 8 cir 20000 pir 50000 cbs
4 ebs 32
```


Copyright© 2012-2013 Ciena® Corporation

8 Create the child profile to classify traffic with DSCP value 16.

```
traffic-profiling standard-profile create port 6 name
child2dscp parent parent2 dscp 16 cir 15000 pir 50000 cbs
4 ebs 16
```

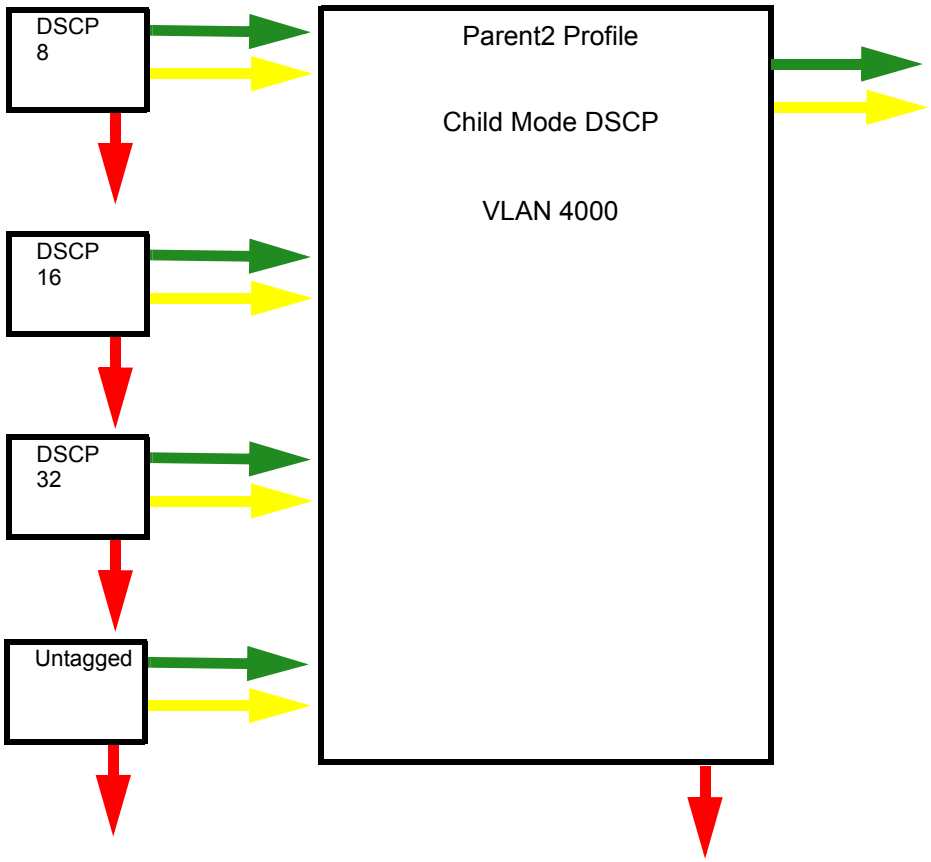
9 Create the child profile to classify traffic with DSCP value 8.

```
traffic-profiling standard-profile create port 6 name
child3dscp parent parent2 dscp 32 cir 14000 pir 50000 cbs
4 ebs 16
```

10 Create the child profile to classify and drop untagged traffic.

```
traffic-profiling standard-profile create port 1 name
child4untagged parent parent2 drop true untagged
```

Figure 13-6
Hierarchical VLAN with DSCP priority



11 Create the parent profile to classify traffic with VLAN 2100 with a narrow classifier-mode and specify the child mode to classify based on IP precedence.

```
traffic-profiling standard-profile create port 6 name
parentIpp child-mode standard-ip-prec cir 50000 cbs 8
vlan 2100
```


Copyright© 2012-2013 Ciena® Corporation

12 Create the child profile to classify traffic with IP precedence value 6.

```
traffic-profiling standard-profile create port 6 name
child1ipp parent parentIpp ip-prec 6 cir 20000 pir 50000
cbs 4 ebs 16
```

13 Create the child profile to classify traffic with IP precedence value 5.

```
traffic-profiling standard-profile create port 6 name
child2ipp parent parentIpp ip-prec 5 cir 10000 pir 50000
cbs 4 ebs 16
```

14 Create the child profile to classify traffic with IP precedence value 4.

```
traffic-profiling standard-profile create port 6 name
child3ipp parent parentIpp ip-prec 4 cir 19000 pir 50000
cbs 4 ebs 16
```

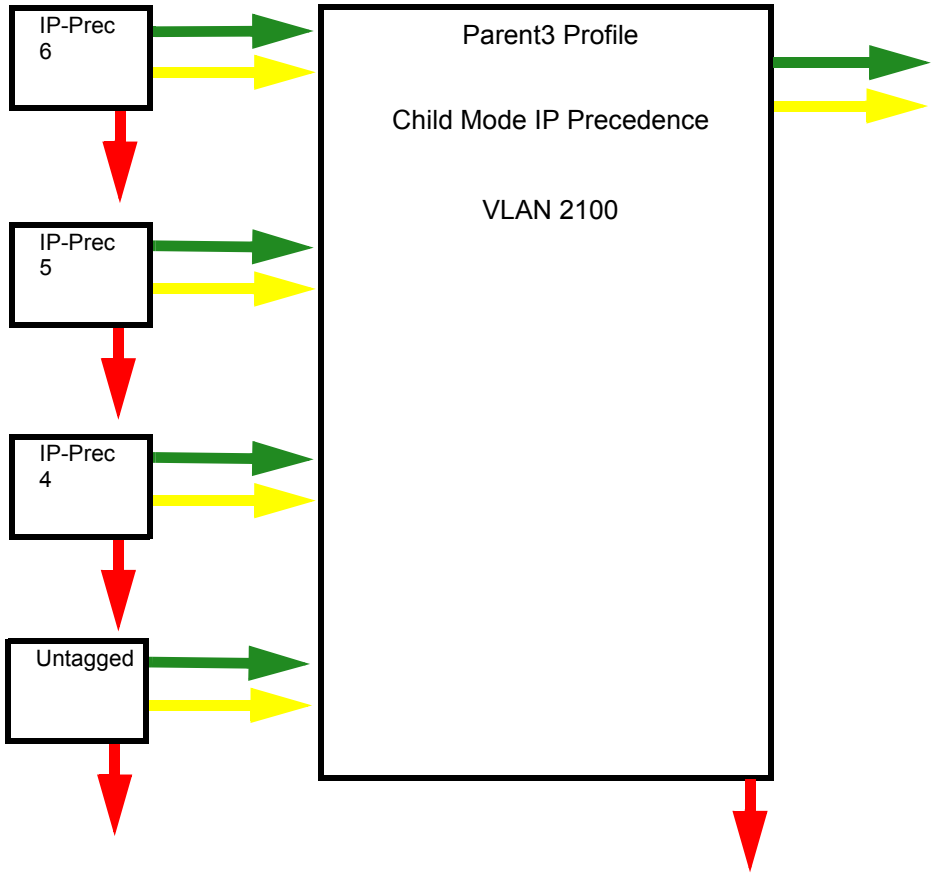
15 Create the child profile to classify and drop untagged traffic.

```
traffic-profiling standard-profile create port 1 name
child4ipp-untagged parent parentIpp drop true untagged
```

Note: Alternately, an profile can be created at the same level as the parent to classify untagged traffic.

Copyright© 2012-2013 Ciena® Corporation

Figure 13-7
Hierarchical VLAN with IP precedence



16 Create a default profile at the parent level for handling traffic that does not classify to the parent and child profiles.

```
traffic-profiling standard-profile create port 6 default-
profile true cir 15000 pir 45000 cbs 8 ebs 8 name
T1default
```

—end—

Copyright© 2012-2013 Ciena® Corporation

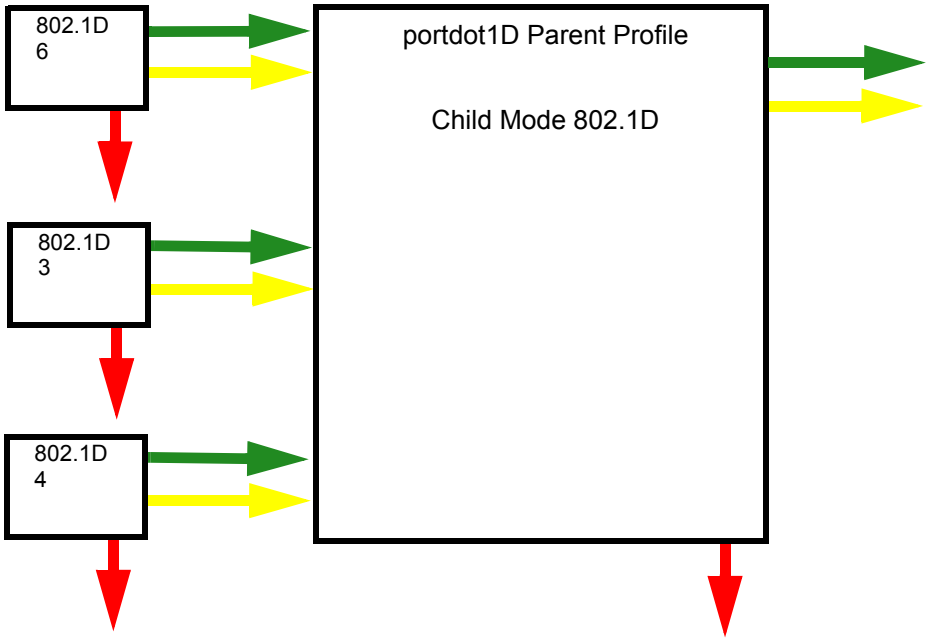
Procedure 13-25
Configuring hierarchical port traffic profiles

In the configuration example for this section, one parent profile with the 802.1D child mode and three child profiles for classification of specific 802.1D values.

Step	Action
1	Set the port mode to hierarchical port mode. `traffic-profiling set port 5 mode hierarchical-port`
2	Create the parent profile with a narrow classifier-mode and the child mode to classify based on 802.1D priority. `traffic-profiling standard-profile create port 5 name portdot1D child-mode standard-dot1dpri cir 50000 cbs 8`
3	Create the child profile to classify traffic with 802.1D priority values 6. `traffic-profiling standard-profile create port 5 name childdot1D1 parent portdot1D dot1dpri 6 cir 30000 pir 40000 cbs 4 ebs 2`
4	Create the child profile to classify traffic with 802.1D priority values 3. `traffic-profiling standard-profile create port 5 name childdot1D2 parent portdot1D dot1dpri 3 cir 10000 pir 25000 cbs 4 ebs 8`
5	Create the child profile to classify traffic with 802.1D priority values 4. `traffic-profiling standard-profile create port 5 name childdot1D3 parent portdot1D dot1dpri 4 cir 9800 pir 9800 cbs 4 ebs 16`

Copyright© 2012-2013 Ciena® Corporation

Figure 13-8
Hierarchical port with 802.1D priority



Copyright© 2012-2013 Ciena® Corporation

Procedure 13-26
Configuring VS classification for standard traffic profiles

In the configuration example for this section, profile 1 classifies on all traffic from port 1 that belongs to VS1 regardless of VLAN or COS. Traffic-profiles 2-4 only classify on traffic ingressing VS2 that have VLAN 2002 and the respective CoS values provisioned for each profile.

Step	Action

1	Create the virtual switches and add them to the respective VLANs.
	`virtual-switch add reserved-vlan 4091-4093`
	`virtual-switch ethernet create vs VS1`
	`virtual-switch Ethernet add vs VS1 port 1 vlan 1001`
	`virtual-switch Ethernet add vs VS1 port 1 vlan 1002`
	`virtual-switch ethernet create vs VS2`
	`virtual-switch Ethernet add vs VS2 port 1 vlan 2001`
	`virtual-switch Ethernet add vs VS2 port 1 vlan 2002`
	`virtual-switch ethernet create vs VS3`
	`virtual-switch Ethernet add vs VS3 port 1 vlan 3001`
	`virtual-switch Ethernet add vs VS3 port 1 vlan 3002`

2	Enable traffic profiling and set the mode to advanced.
	`traffic-profiling enable`
	`traffic-profiling enable port 1`
	`traffic-profiling set port 1 mode advanced`

3	Create traffic profile 1 and assign it to VS1, specifying CIR and PIR values. `traffic-profiling standard-profile create port 1 profile 1 cir 1024 pir 2048 vs VS1`

4	Create traffic profiles 2-4 and assign them to VS2, specifying CIR, PIR, VLAN, and 802.1D values.
	`traffic-profiling standard-profile create port 1 profile 2 cir 4096 pir 8192 vs VS2 vlan 2002 dot1dpri 0,1,2,3`
	`traffic-profiling standard-profile create port 1 profile 3 cir 512 pir 1024 vs VS2 vlan 2002 dot1dpri 4,5`
	`traffic-profiling standard-profile create port 1 profile 4 cir 2048 pir 2048 vs VS2 vlan 2002 dot1dpri 6,7`

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-27
Configuring VS classification for HIM traffic profiles

In the configuration example for this section, regardless of the VLAN, any traffic ingressing VS3 with the specified CoS values is classified/metered.

Step	Action
1	Create the virtual switch configuration as described in Procedure 13-26.

2	Enable traffic profiling and set the mode to advanced.
	`traffic-profiling enable`
	`traffic-profiling enable port 1`
	`traffic-profiling set port 1 mode advanced`

3	Create parent traffic profile Parent1 and assign it to VS3 with child-mode VLAN-CoS, specifying CIR and PIR values. `traffic-profiling standard-profile create port 1 name Parent1 child-mode vlan-cos cir 50048 vs VS3`

4	Create child traffic profiles Child1 and Child2 and assign them to profile Parent1, specifying CIR, PIR, and DSCP values.
	`traffic-profiling standard-profile create port 1 name Child1 parent Parent1 cir 1024 pir 2048 dscp 0,1,2,3`
	`traffic-profiling standard-profile create port 1 name Child2 parent Parent1 cir 1024 pir 2048 dscp 4,5`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-28
Displaying standard traffic profiles

You can display

- all standard traffic profiles
- standard traffic profiles associated with a specified port

Step	Action

To display all standard traffic profiles

1 Display all standard traffic profiles:

```
traffic-profiling standard-profile show
```

To display standard traffic profiles associated with a specified port

2 Display standard traffic profiles associated with a specified port:

```
traffic-profiling standard-profile show [port
<PortName>]>
```

where

<PortName> is the port that you want to display associated standard traffic profiles for.

—end—

Example

The folliowing example shows sample output for displaying all standard traffic profiles.

```
traffic-profiling standard-profile show
```

```
+-----------------------------------+- STANDARD PROFILE TABLE-------------------------------+
| Port  |      Profile      |     |Parent|    BW (Kbps)   | Max Burst KB    |    CLASSIFIERS   |
|       | ID| Name          |Role|#Child| CIR   | PIR   |   CBS  |  EBS   |                  |
+-------+---+---------------+----+------+-------+-------+--------+--------+------------------+
| 1     |1  |STD#1          | N  |   0  |128    |128    | 16     | 16     | .1D| 1           |
|       |   |               |    |      |       |       |        |        | IPP| 1           |
+-------+---+---------------+----+------+-------+-------+--------+--------+------------------+
| 2     |1  |STD#1          | N  |   0  |0      |256    | 0      | 16     |                  |
+-------+---+---------------+----+------+-------+-------+--------+--------+------- ----------+
| 2     |2  |untagged       | N  |   0  | 128   |256    | 16     | 16     | UT               |
+-------+---+---------------+----+------+-------+-------+--------+--------+------------------+
```

The following example shows sample output for displaying standard profiles associated with port 1.

```
traffic-profiling standard-profile show port 1
```

```
+-----------------------------------+- STANDARD PROFILE TABLE-------------------------------+
| Port  |      Profile      |     |Parent|    BW (Kbps)   | Max Burst KB    |    CLASSIFIERS   |
|       | ID| Name          |Role|#Child| CIR   | PIR   |   CBS  |  EBS   |                  |
```


Copyright© 2012-2013 Ciena® Corporation

```
+-------+---+---------------+----+------+-------+-------+--------+--------+------------------+
| 1     |1  |STD#1          | N  |   0  |128    |128    | 16     | 16     | .1D| 1           |
|                                                                          | IPP| 1           |
+-------+---+---------------+----+------+-------+-------+--------+--------+------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-29
Clearing statistics for all standard profiles

Statistics count the number of accepted and dropped bytes and packets for each standard profile.

Note 1: Egress and traffic profile byte statistics show “payload” bytes regardless of the frame bandwidth calculation setting described in the “Configuring frame bandwidth calculation” procedure.

Note 2: The 3902, 3911, 3920, 3940, and 5140 platforms only support statistics for dropped bytes and accepted bytes.

Step	Action
1	Clear statistics for all standard profiles: `traffic-profiling standard-profile clear statistics`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-30
Displaying statistics for all standard profiles

Display statistics for all standard profiles.

Step	Action
1	Display statistics for all standard profiles: `traffic-profiling standard-profile show statistics`

—end—

Example

The following example shows sample output for displaying statistics for all standard profiles.

```
traffic-profiling standard-profile show statistics

+------------------------------- STANDARD PROFILE TABLE ---------------------------------+
|  Port |      Profile             | Statistics                                          |
|       | ID | Name                | Type       Bytes                  packets           |
+-------+----+---------------------+-----------------------------------------------------+
|1      |1   |p1                   | **** Child Profile List ****                        |
+-------+----+---------------------+-------------------------------+---------------------+
|1      |2   |     child1          | Accepted   0                  | 0                   |
|       |    |                     | Dropped    0                  | 0                   |
+-------+----+---------------------+-------------------------------+---------------------+
|       |    |                     | **** Child Profiles Sum ****                        |
+-------+----+---------------------+-------------------------------+---------------------+
|1      |1   |p1                   | Accepted   0                  | 0                   |
|       |    |                     | Dropped    0                  | 0                   |
+-------+----+---------------------+-------------------------------+---------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-31
Displaying throughput statistics for a traffic profile

Display throughput statisitcs for a traffic profile to monitor the troughput traffic profile from the CLI. Throughput statistics are used for troubleshooting.

Step	Action
1	Display throughput statistics for a traffic profile: `traffic profiling standard-profile show throughput port <PortName> traffic-profile <TrafficProfileName>` where port <PortName> is the port that the traffic profile resides on traffic-profile <TrafficProfile Name> is the traffic profile that you want to monitor

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-32
Setting the per-port hierarchical traffic-profiling mode

Set the per-port hierarchical traffic-profiling mode.

Step	Action
1	Set the per-port hierarchical traffic-profiling mode: `traffic-profiling set port <PortName> mode <none|advanced|standard-dot1dpri|standard-ip-prec|standard-dscp|standard-vlan|standard-vlan-dot1dpri|standard-vlan-ip-prec|standard-vlan-dscp|hierarchical-port|hierarchical-vlan>`

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-33
Setting the child mode

Set the child mode (child-mode) attribute to determine the classification of child traffic profiles:

- standard-dot1dpri - Finds a matching traffic profile using 802.1D priority. This mode is the default.
- standard-ip-prec - Finds a matching profiling using the upper 3 bits of the TOS byte that make up the IP precedence.
- standard-dscp - Finds a matching profiling using the DSCP value.
- standard-vlan - Finds a matching profiling based upon the VLAN ID.
- vlan-cos - Finds a matching profiling based upon the VLAN CoS. This child mode is paired with a parent level EVC(VS) classification and it allows for child profiles to have a VLAN, a VLAN + CoS, or CoS only, as long as the child profiles don't overlap classification.

  Note: If the parent profile is using a VS classifier, `vlan-cos` is the only acceptable child mode, and the traffic profiling mode for the parent port must be set to `advanced`.

Step	Action
1	Set the child mode: `traffic-profiling standard-profile set {port <PortNameList>} {profile <TrafficProfStd>} child-mode <standard-dot1dpri|standard-ip-prec|standard-dscp|standard-vlan|vlan-cos>` **—end—**

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-34
Creating an sRED profile

Create an sRED profile.

Step	Action
1	Create an sRED profile: `traffic-services queuing congestion-avoidance-profile create profile <String[15]> [green-threshold <NUMBER: 1-100>] [green-drop-probability <100pct|6.25pct|3.125pct|1.5625pct|0.78125pct|0.390625pct|0.1953125pct|0.09765625pct>] [yellow-threshold <NUMBER: 1-100>] [yellow-drop-probability <100pct|6.25pct|3.125pct|1.5625pct|0.78125pct|0.390625pct|0.1953125pct|0.09765625pct>]]`

—end—

Example

The following example creates an sRED profile.

```
traffic-services queuing congestion-avoidance-profile create profile green-yellow-1 green-threshold 100 yellow-threshold 1 yellow-drop-probability 100pct
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-35
Modifying an sRED profile

Modify an sRED profile.

Step	Action
1	Modify an sRED profile: `traffic-services queuing congestion-avoidance-profile set profile <CongestionAvoidanceProfile> {[green-threshold <NUMBER: 1-100>], [green-drop-probability <100pct|6.25pct|3.125pct|1.5625pct|0.78125pct|0.390625pct|0.1953125pct|0.09765625pct>], [yellow-threshold <NUMBER: 1-100>], [yellow-drop-probability <100pct|6.25pct|3.125pct|1.5625pct|0.78125pct|0.390625pct|0.1953125pct|0.09765625pct>]}` **—end—**

Example

The following example modifies an sRED profile.

```
traffic-services queuing congestion-avoidance-profile set profile green-
yellow-1 green-threshold 90 green-drop-probability 0.1953125pct yellow-
threshold 25
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-36
Creating an sWRED profile

Create an sWRED profile.

Step	Action
1	Create an sWRED profile: `traffic-services queuing congestion-avoidance-profile create profile <String[15]> [tcp-green-threshold <NUMBER: 1-100>] [tcp-green-upper-threshold <NUMBER: 1-100>] [tcp-green-max-drop-probability <100pct|75pct|50pct|25pct|10pct|9pct|8pct|7pct|6pct|5pct|4pct|3pct|2pct|1pct|0pct>] [tcp-yellow-threshold <NUMBER: 1-100>] [tcp-yellow-upper-threshold <NUMBER: 1-100>] [tcp-yellow-max-drop-probability <100pct|75pct|50pct|25pct|10pct|9pct|8pct|7pct|6pct|5pct|4pct|3pct|2pct|1pct|0pct>] [non-tcp-threshold <NUMBER: 1-100>] [non-tcp-upper-threshold <NUMBER: 1-100>] [non-tcp-max-drop-probability <100pct|75pct|50pct|25pct|10pct|9pct|8pct|7pct|6pct|5pct|4pct|3pct|2pct|1pct|0pct>]` **—end—**

Example

The following example creates an sWRED profile.

```
traffic-services queuing congestion-avoidance-profile create profile green-yellow-1 tcp-green-upper-threshold 100 tcp-yellow-upper-threshold 1 tcp-yellow-max-drop-probability 100pct
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-37
Modifying an sWRED profile

Modify an sWRED profile.

Step	Action
1	Modify an sWRED profile: `traffic-services queuing congestion-avoidance-profile set profile <CongestionAvoidanceProfile> [tcp-green-threshold <NUMBER: 1-100>] [tcp-green-upper-threshold <NUMBER: 1-100>] [tcp-green-max-drop-probability <100pct|75pct|50pct|25pct|10pct|9pct|8pct|7pct|6pct|5pct|4pct|3pct|2pct|1pct|0pct>] [tcp-yellow-threshold <NUMBER: 1-100>] [tcp-yellow-upper-threshold <NUMBER: 1-100>] [tcp-yellow-max-drop-probability <100pct|75pct|50pct|25pct|10pct|9pct|8pct|7pct|6pct|5pct|4pct|3pct|2pct|1pct|0pct>] [non-tcp-threshold <NUMBER: 1-100>] [non-tcp-upper-threshold <NUMBER: 1-100>][non-tcp-max-drop-probability <100pct|75pct|50pct|25pct|10pct|9pct|8pct|7pct|6pct|5pct|4pct|3pct|2pct|1pct|0pct>]` **—end—**

Example

The following example modifies an sWRED profile.

```
traffic-services queuing congestion-avoidance-profile set profile green-
yellow-1 non-tcp-upper-threshold 1 non-tcp-max-drop-probability 100pct
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-38
Creating an sWRED profile

Create an sWRED profile.

Step	Action
1	Create an sWRED profile: `traffic-services queuing congestion-avoidance-profile create profile <String[15]> [green-lower-threshold <NUMBER: 1-100>] [green-upper-threshold <NUMBER: 1-100>] [green-max-drop-probability <100pct|75pct|50pct|25pct|10pct|9pct|8pct|7pct|6pct|5pct|4pct|3pct|2pct|1pct|0pct>] [yellow-lower-threshold <NUMBER: 1-100>] [yellow-upper-threshold <NUMBER: 1-100>] [yellow-max-drop-probability <100pct|75pct|50pct|25pct|10pct|9pct|8pct|7pct|6pct|5pct|4pct|3pct|2pct|1pct|0pct>]` **—end—**

Example

The following example creates an sWRED profile.

```
traffic-services queuing congestion-avoidance-profile create profile green-yellow-1 green-upper-threshold 100 yellow-upper-threshold 1 yellow-max-drop-probability 100pct
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-39
Modifying an sWRED profile

Modify an sWRED profile.

Step	Action
1	Modify an sWRED profile: `traffic-services queuing congestion-avoidance-profile set profile <CongestionAvoidanceProfile> [green-lower-threshold <NUMBER: 1-100>] [green-upper-threshold <NUMBER: 1-100>] [green-max-drop-probability <100pct|75pct|50pct|25pct| 10pct|9pct|8pct|7pct|6pct|5pct|4pct|3pct|2pct|1pct|0pct>] [yellow-lower-threshold <NUMBER: 1-100>] [yellow-upper-threshold <NUMBER: 1-100>] [yellow-max-drop-probability <100pct|75pct|50pct|25pct|10pct|9pct|8pct|7pct|6pct|5pct |4pct|3pct| 2pct|1pct|0pct>]` **—end—**

Example

The following example modifies an sWRED profile.

```
traffic-services queuing congestion-avoidance-profile set profile green-
yellow-1 yellow-upper-threshold 50 yellow-max-drop-probability 100pct
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-40
Displaying custom congestion avoidance profiles

You can

- display all congestion avoidance profiles
- display specific congestion avoidance profiles, including name, ID, type, thresholds, and drop probability

Step	Action

To display all congestion avoidance profiles

1 Display all congestion avoidance profiles:

```
traffic-services queuing congestion-avoidance-profile
show
```

To display a specific congestion avoidance profiles

2 Display a specific congestion avoidance profile:

```
traffic-services queuing congestion-avoidance-profile
show profile <CongestionAvoidanceProfile>
```

where

profile <CongestionAvoidanceProfile>	identifies the specific congestion avaoidance profile.

—end—

Examples

The following example shows the default sRED congestion avoidance profile.

```
traffic-services queuing congestion-avoidance-profile show profile
Default-SRED

+------------------ CONGESTION AVOIDANCE PROFILE DATA --------------------+
|                                                                          |
+------------------------------+-------------------------------------------+
| Name                         | Default-SRED                              |
| Id                           | 1                                         |
| Type                         | SRED                                      |
+------------------------------------------------------------------------+
|          Green                          ||          Yellow              |
| Threshold |    Drop Probability         || Threshold |    Drop Probability  |
+-----------+-------------------------++-----------+-------------------------+
| 75      % |0.09765625pct.(1/1024    || 50      % | 6.25pct.......(1/16)    |
+-----------+-------------------------++-----------+-------------------------+
```

The following example shows the default sWRED congestion avoidance profile on the 3960.

Copyright© 2012-2013 Ciena® Corporation

```
traffic-services queuing congestion-avoidance-profile show profile Default-S-
WRED

+------------------------------ CONGESTION AVOIDANCE PROFILE DATA ------------------------------+
|                                                                                               |
+-------------------------------+-----------------------------------------------------------------+
| Name                          | Default-S-WRED                                                  |
| Id                            | 1                                                               |
| Type                          | WRED-Simple                                                     |
+-------------------------------+-----------------------------------------------------------------+
|             Tcp-Green         ||            Tcp-Yellow         ||   Non-Tcp                    |
|  Lower    | Upper     | Drop  ||  Lower    | Upper     | Drop  ||  Lower    | Upper     | Drop  |
| Threshold | Threshold | Prob  || Threshold | Threshold | Prob  || Threshold | Threshold | Prob  |
+-----------+-----------+-------++-----------+-----------+-------++-----------+-----------+-------+
| 75        | 100       | 0pct  || 50        | 100       | 100pct || 75        | 100       | 0pct  |
+-----------+-----------+-------++-----------+-----------+-------++-----------+-----------+-------+
```

The following example shows the default sWRED congestion avoidance profile on the 5150.

```
traffic-services queuing congestion-avoidance-profile show profile Default-S-
WRED

traffic-services queuing congestion-avoidance-profile show

+--------------- CONGESTION AVOIDANCE PROFILE DATA ----------------+
|                                                                  |
+---------------------------------+--------------------------------+
| Name                            | Default-S-WRED                 |
| Id                              | 1                              |
| Type                            | WRED-Simple                    |
+---------------------------------+--------------------------------+
|               Green             ||              Yellow            |
|  Lower    | Upper     | Drop    ||  Lower    | Upper     | Drop   |
| Threshold | Threshold | Prob    || Threshold | Threshold | Prob   |
+-----------+-----------+---------++-----------+-----------+--------+
| 75     %  | 100    %  | 0pct    || 50     %  | 100    %  | 100pct |
|                                                                  |
+---------------------------------+--------------------------------+
| Name                            | custom-default                 |
| Id                              | 2                              |
| Type                            | WRED-Simple                    |
+---------------------------------+--------------------------------+
|               Green             ||              Yellow            |
|  Lower    | Upper     | Drop    ||  Lower    | Upper     | Drop   |
| Threshold | Threshold | Prob    || Threshold | Threshold | Prob   |
+-----------+-----------+---------++-----------+-----------+--------+
| 75     %  | 100    %  | 0pct    || 50     %  | 100    %  | 100pct |
+-----------+-----------+---------++-----------+-----------+--------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-41
Updating the congestion avoidance profile for an egress port queue

The default congestion avoidance profile for egress port queues is to use Default-SRED or Default-S-WRED, depending upon the platform. After creating custom congestion avoidance profiles, you can set the egress port queues to use them. The queue is identified by queue number and egress port queue group (PortQueueGroup), which is the egress port name.

Step	Action
1	Update the congestion avoidance profile for an egress port queue: `traffic-services queuing egress-port-queue-group set queue <NUMBER: 0-7> port <PortQueueGroup> congestion-avoidance-profile <CongestionAvoidanceProfileName>`

—end—

Example

The following example updates the congestion avoidance profile for an egress port queue.

```
traffic-services queuing egress-port-queue-group set queue 2 port 1
congestion-avoidance-profile green-yellow-1
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-42
Clearing the congestion avoidance profile to the default for an egress port queue

Clear the congestion avoidance profile to the default for an egress port queue.

Step	Action
1	Clear the congestion avoidance profile to the default for an egress port queue: `traffic-services queuing egress-port-queue-group unset queue <NUMBER: 0-7> port <PortQueueGroup> congestion-avoidance-profile` **—end—**

Example

The following example clears the congestion avoidance profile to the default.

```
traffic-services queuing egress-port-queue-group unset queue 2 port 1
congestion-avoidance-profile
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-43
Displaying the congestion avoidance profile for an egress port queue

Display the congestion avoidance profile for an egress port queue.

Step	Action
1	To display the congestion avoidance profile for an egress port queue: `traffic-services queuing egress-port-queue-group show queue <NUMBER: 0-7> port <PortQueueGroup>`

—end—

Example

Example of sRED:

```
traffic-services queuing egress-port-queue-group show queue 0 port 1

+------------------------------- QUEUE DATA -----------------------------+
|                                                                        |
+--------------------------------+---------------------------------------+
| Queue Id                       | 0                                     |
| Queue Group Name [Port]        | 1                                     |
| Congestion Avoidance Profile   | Default-SRED                          |
+--------------------------------+---------------------------------------+
|     Scheduler      |   Size    |    CIR     |    CBS     |    EIR     |    EBS     |
| Pri Idx | Weight   | (Pckts)   |   (Kbps)   | (Kbytes)   |   (Kbps)   | (Kbytes)   |
+---------+----------+-----------+------------+------------+------------+------------+
| 10000   | 20       | 100       |  0         | 0          |1000000     | 256        |
+---------+----------+-----------+------------+------------+------------+------------+
```

Example of sWRED on the 3960:

```
traffic-services queuing egress-port-queue-group show queue 0 port 1

+------------------------------- QUEUE DATA -----------------------------+
|                                                                        |
+--------------------------------+---------------------------------------+
| Queue Id                       | 0                                     |
| Queue Group Name [Port]        | 1                                     |
| Congestion Avoidance Profile   | Default-S-WRED                        |
+--------------------------------+---------------------------------------+
|     Scheduler      |   Size    |    CIR     |    CBS     |    EIR     |    EBS     |
| Pri Idx | Weight   | (Pckts)   |   (Kbps)   | (Kbytes)   |   (Kbps)   | (Kbytes)   |
+---------+----------+-----------+------------+------------+------------+------------+
| 10000   | 20       | 100       |  0         | 0          |1000000     | 256        |
+---------+----------+-----------+------------+------------+------------+------------+
```

Example of sWRED on the 5150:

```
traffic-services queuing egress-port-queue-group show queue 0 port 1.1
```


Copyright© 2012-2013 Ciena® Corporation

```
+------------------------------- QUEUE DATA -------------------------------+
|                                                                          |
+--------------------------------+-----------------------------------------+
| Queue Id                       | 0                                       |
| Queue Group Name [Port]        | 1.1                                     |
| Congestion Avoidance Profile   | Default-S-WRED                          |
+--------------------------------------------------------------------------+
|     Scheduler      |  Size   |   CIR    |   CBS    |   EIR    |   EBS    |
| Pri Idx | Weight   | (Pckts) |  (Kbps)  | (Kbytes) |  (Kbps)  | (Kbytes) |
+---------+----------+---------+----------+----------+----------+----------+
| 10000   | 20       | 100     |  0       | 0        |1000000   | 256      |
+---------+----------+---------+----------+----------+----------+----------+
```

Procedure 13-44
Deleting a custom congestion avoidance profile

Delete custom congestion avoidance profiles that are not in use.

Step	Action
1	Delete a congestion avoidance profile: `traffic-services queuing congestion-avoidance-profile delete profile <CongestionAvoidanceProfile>` where profile <CongestionAvoidanceProfile> identifies the congestion avaoidance profile to be deleted.

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-45
Renaming custom congestion avoidance profile

You can change the name of a custom congestion avoidance profile.

Step	Action
1	Rename a congestion avoidance profile: `traffic-services queuing congestion-avoidance-profile rename profile <CongestionAvoidanceProfile> name <String[15]>` where profile <CongestionAvoidanceProfile> identifies the congestion avaoidance profile to be renamed. name <String[15]> is the new name for the congestion avaoidance profile. **—end—**

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-46
Changing the algotrithm of an egress port scheduler

Change the algorithm of an egress port scheduler.

Note: The wdrr-scheduler-granularity setting can be set with scheduler algorithms other than WDRR, but does not go into effect until the scheduler algorithm is changed to WDRR. The default setting is 1600 Kb.

Step	Action
1	Change an egress port's scheduler algorithm and granularity for the egress port queue group: `traffic-services queuing egress-port-queue-group set port <PortQueueGroup> {[scheduler-algorithm <strict|round-robin|weighted-deficit-round-robin|weighted-round-robin>], [wdrr-scheduler-granularity <NUMBER: 50-1600>]}` **—end—**

Example

The following example changes an egress port's scheduler algorithm and granularity for the egress port queue group.

```
traffic-services queuing egress-port-queue-group set port 9 scheduler-
algorithm weighted-round-robin
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-47
Changing the weight of the scheduler for a queue

Change the weight of the scheduler for a queue.

Note: The default weight for CoS Queue 7 is 0, which is serviced in a strict-priority fashion being the highest priority over all other CoS queues in the Egress Queue group. Queue 7 is typically reserved for CPU-Sourced traffic and should not be changed.

Step	Action
1	Change a queue's scheduler weight for use with the port's scheduler algorithm: `traffic-services queuing egress-port-queue-group set queue <NUMBER: 0-7> port <PortQueueGroup> scheduler-weight <NUMBER: 0-1000>`

—end—

Example

The following example changes a queue's scheduler weight.

```
traffic-services queuing egress-port-queue-group set queue 0 port 9 scheduler-weight 100
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-48
Displaying queue weight and scheduler algorithms

Display queue weight and scheduler algorithms.

Step	Action
1	Display queue weight and scheduler algorithms for all or a specific egress port: `traffic-services queuing egress-port-queue-group show [port <PortQueueGroup>]` where port <PortQueueGroup> is the specific egress port.

—end—

Example

The following example displays queue weight and scheduler algorithms.

```
traffic-services queuing egress-port-queue-group show port 9

+---------------------------- QUEUE GROUP DATA ----------------------------+
|                                                                          |
+----------------------------------+-----------------------------------------+
| Name [Port]                      | 9                                       |
| Id                               | 9                                       |
| Queue Count                      | 8                                       |
| Scheduling Algorithm             | weighted-round-robin                    |
| Shaper Bandwidth (Kbps)          | 1000000                                 |
| Shaper Burst Size (Kb)           | 10240                                   |
+----------------------------------------------------------------------------+
|   |      Scheduler      |   Size    |    CIR    |    CBS    |    EIR    |    EBS    |
| Q | Pri Idx | Weight    | (Pckts)   | (Kbps)    | (Kbytes)  | (Kbps)    | (Kbytes)  |
+---+---------+-----------+-----------+-----------+-----------+-----------+-----------+
| 0 | 10000   | 100       | 100       | 0         | 0         | 1000000   | 256       |
| 1 | 20000   | 30        | 100       | 0         | 0         | 1000000   | 256       |
| 2 | 30000   | 40        | 100       | 0         | 0         | 1000000   | 256       |
| 3 | 40000   | 50        | 100       | 0         | 0         | 1000000   | 256       |
| 4 | 50000   | 60        | 100       | 0         | 0         | 1000000   | 256       |
| 5 | 60000   | 70        | 100       | 0         | 0         | 1000000   | 256       |
| 6 | 70000   | 80        | 100       | 0         | 0         | 1000000   | 256       |
| 7 | 80000   | 0         | 100       | 102       | 256       | 1000000   | 256       |
+---+---------+-----------+-----------+-----------+-----------+-----------+-----------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-49
Configuring egress port and queue shaping

You can set

- egress port shaper-rate bandwidth
- egress port shaper-rate burst-size
- egress shaping parameter for each queue

Table 13-8 lists the steps that the shaper rate increments in based on the burst size.

Table 13-8
Burst size to shaper-rate increment

Burst size	Shaper rate increments in steps of
0 or less than 16 M	64 kbps ***Note:*** 3902, 3911, 3920, 3940, 5140 switches do not exceed 16 M, and always use this increment.
16 M and less than 32 M	128 kbps
32 M and less than 64 M	256 kbps
64 M and less than 128 M	512 kbps
128 M and less than 512 M	1024 kbps

Table 13-9 lists the steps that the CIR and PIR increments in based on the shaper rate.

Table 13-9
Shaper rate to CIR and PIR increment

Shaper rate	CIR and PIR increments in steps of
0 or less than 16.78 Gbps	64 kbps ***Note:*** 3902, 3911, 3920, 3940, 5140 switches do not exceed 16 M, and always use this increment.
16.78 Gbps and less than 33.55 Gbps	128 kbps
33.55 Gbps and less than 40 Gbps	256 kbps

Copyright© 2012-2013 Ciena® Corporation

Table 13-10 lists the parameters for individual shaping of queues.

Table 13-10
Parameters for individual shaping of queues

Parameter	Description
CIR in Kbps	The rate of traffic that can egress from a queue and be considered guaranteed traffic. The default CIR for queues 0-6 is 0, queue 7 is 1024. The CIR value ranges from 0 to 40 Gbps (40000000 Kbps).
CBS in Kbytes	The amount of CIR traffic that can burst from a queue (the CIR bucket size). The default CBS for queues 0-6 is 0, queue 7 is 256. The CBS value ranges from 0 to 256 Mbytes.
EIR in Kbps	The rate of traffic that can egress from a queue above CIR and is considered non-guaranteed traffic. The default EIR is the same as the administrative port speed. The EIR value ranges from 0 to 40 Gbps.
EBS in Kbytes	The amount of EIR traffic that can egress from a queue (the EIR bucket size). The default EBS is 256. The EBS value ranges from 0 to 256 Mbytes.

Note 1: In order for the EBS bucket to be drained, an EIR value must be set.

Note 2: For 3902, 3911, 3920, 3940, 5140 switches the sum of the CBS and EBS can be less than or equal to 16 Mbytes. For 3916, 3930, 3931, 3960, 5150 switches, the sum of the CBS and EBS can be less than or equal to 256 Mbytes

Note 3: Shaper-rates applied at the port level affect all 8 queues in the egress queue group.

Step	Action

To set the egress port shaper-rate bandwidth

1 Set the egress port shaper-rate bandwidth:

```
traffic-services queuing egress-port-queue-group set port
<PortQueueGroup> shaper-rate <NUMBER: 0-40000000>
```

To set the egress port shaper-rate burst-size

2 Set the egress port shaper-rate burst-size:

```
traffic-services queuing egress-port-queue-group set port
<PortQueueGroup> burst-size <NUMBER: 4-10240>
```


Copyright© 2012-2013 Ciena® Corporation

To set the egress shaping parameters for each queue

3 Set the egress shaping parameters for each queue:

```
traffic-services queuing egress-port-queue-group set
queue <NUMBER: 0-7> port <PortQueueGroup> cir <NUMBER: 0-
20000000> cbs <NUMBER: 0-512> eir <NUMBER: 0-20000000>
ebs <NUMBER: 0-512>
```

—end—

Example

The following example sets the egress port shaper-rate bandwidth.

```
traffic-services queuing egress-port-queue-group set port 3 shaper-rate 1200
```

The following example sets the egress port shaper-rate burst-size.

```
traffic-services queuing egress-port-queue-group set port 3 burst-size 1024
```

The following example sets the egress shaping parameter per queue.

```
traffic-services queuing egress-port-queue-group set queue 0 port 9 cir 5000
cbs 5000 eir 4000 ebs 3000
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 13-50
Displaying egress port queue configuration

You can display egress port queue configuration for

- all ports
- a queue group on a port
- a port

Note: Egress and traffic profile byte statistics show “payload” bytes regardless of the frame bandwidth calculation setting described in the “Configuring frame bandwidth calculation” procedure.

Step	Action

To display egress port queue configuration for all ports

1 Display egress port queue configuration for all ports:

```
traffic-services queuing egress-port-queue-group show
```

where

queue <NUMBER: 0-7> — is the queue in the queue group

port <PortQueueGroup> — is the port on which the queue group is instantiated.

To display egress port queue configuration for a queue group on a port

2 Display egress port queue configuration for a queue group on a port:

```
traffic-services queuing egress-port-queue-group show
queue <NUMBER: 0-7> port <PortQueueGroup>
```

To display egress port queue configuration for a port

3 Display egress port queue configuration for a port:

```
traffic-services queuing egress-port-queue-group show
port <PortQueueGroup>
```

—end—

Example

```
aggregation create aggregation agg1

traffic-services queuing egress-port-queue-group show port agg1
+---------------------------- QUEUE GROUP DATA ----------------------------+
|                                                                          |
+---------------------------------+----------------------------------------+
| Name [Port]                     | agg1                                   |
| Id                              | 2049                                   |
| Queue Count                     | 8                                      |
```


Copyright© 2012-2013 Ciena® Corporation

Scheduling Algorithm	strict
Shaper Bandwidth (Kbps)	20000000
Shaper Burst Size (Kb)	10240

Q	Scheduler Pri	Scheduler Idx	Scheduler Weight	Size (Pckts)	CIR (Kbps)	CBS (Kbytes)	EIR (Kbps)	EBS (Kbytes)
0	10000	20	100	0	0	20000000	256	
1	20000	30	100	0	0	20000000	256	
2	30000	40	100	0	0	20000000	256	
3	40000	50	100	0	0	20000000	256	
4	50000	60	100	0	0	20000000	256	
5	60000	70	100	0	0	20000000	256	
6	70000	80	100	0	0	20000000	256	
7	80000	0	100	1024	256	20000000	256	

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-51
Displaying egress port queue statistics

Statistics count the bytes and packets dropped and transmitted by the egress port queues.

You can display egress port queue configuration for

- all ports
- a queue group on a port
- a port

Note: Egress and traffic profile byte statistics show “payload” bytes regardless of the frame bandwidth calculation setting described in the “Configuring frame bandwidth calculation” procedure.

Step	Action

To display egress port queue statistics for all ports

1 Display egress port queue configuration for all ports:

```
traffic-services queuing egress-port-queue-group show
statistics
```

where

queue <NUMBER: 0-7>	is the queue in the queue group
port <PortQueueGroup>	is the port on which the queue group is instantiated.

To display egress port queue statistics for a queue group on a port

2 Display egress port queue configuration for a queue group on a port:

```
traffic-services queuing egress-port-queue-group show
queue <NUMBER: 0-7> port <PortQueueGroup> statistics
```

To display egress port queue statistics for a port

3 Display egress port queue configuration for a port:

```
traffic-services queuing egress-port-queue-group show
port <PortQueueGroup> statistics
```

—end—

Copyright© 2012-2013 Ciena® Corporation

Procedure 13-52
Configuring frame bandwidth calculation

The system supports 20 bytes of IFG that can be considered when calculating bandwidth for broadcast containment filters and for traffic profiling according to CIR and PIR settings in a traffic profile for ingress metering and egress shaping.

Step	Action
1	To set the bandwidth calculation mode: `flow bw-calculation-mode set mode <payload | transport>`
2	To NOT use IFG in bandwidth calculations: `flow bw-calculation set mode payload`
3	To use IFG in bandwidth calculations: `flow bw-calculation set mode transport`

Note that when the system is in payload calculation mode, the 20-byte IFG is not considered in the bandwidth equations. However, the physical port must still operate according to Ethernet standards and take into account the IFG when transmitting packets. Therefore, the operator could possibly configure a CIR sum at a port that would exceed the physical bandwidth of the port due to the required IFG on the wire. The software does not attempt to adjust bandwidth or issue user-events when in payload mode. The software assumes bandwidth to be configured and consumed according to transport mode.

Bandwidth rates are calculated differently for Layer 1 and Layer 2 packets as shown in the following:

- R1 - Layer 1 rate
- R2 - Layer 2 rate
- F - Packet frame size

Rate calculation formulas are:

- R1 = R2/(F+20) * F
- R2 = R1/F * (F+20)

Copyright© 2012-2013 Ciena® Corporation

Copyright© 2012-2013 Ciena® Corporation

Multicast services configuration

This chapter provides information on how to configure multicast services in some typical situations, including:

- "IGMP snooping" on page 14-1
- "Unknown multicast filtering" on page 14-8
- "Configuring multicast" on page 14-9

Note: To configure Multicast Services, you need to install the Advanced Ethernet license key. To obtain the Advanced Ethernet license key, contact Ciena Sales.

IGMP snooping

While traditional IGMP snooping works well in most enterprise networks, service providers delivering converged services need to deliver high-bandwidth multicast services to many subscribers without disruptions. These types of networks must handle near-real-time join/leave processing for all subscribers.

In order to support these requirements, several enhancements to traditional IGMP are provided, including Proxy Query Delay, IGMP message shaping, an IGMP query engine, IGMP message filtering, Channel Streams, configurable IGMP leave modes, rapid topology response, and fast join/leave processing. It should be noted that multicast services (IGMP) is globally enabled for the entire device by default, however, it can be enabled or disabled for each multicast VLAN individually.

Note: IGMP is not supported when the remote interface management access is provided over an MPLS VS.

IGMP message shaping

Traditional IGMP snooping sends a proxy query to all devices at once. In response all ports in a device tend to reply at nearly the same time. This causes bursts or spikes of IGMP join messages that all reach the router at the same time. To alleviate this problem, a type of proxy reporting called IGMP

Copyright© 2012-2013 Ciena® Corporation

message shaping is used. When an IGMP query is received, IGMP membership reports are sent at a constant rate to the router, based on the list of active multicast groups. This eases the burden on the multicast querier.

Proxy query delay

Another problem with traditional IGMP snooping occurs on the downstream devices. When a general query is received, it is forwarded to all ports on the device at the same time. This request requires it to update its multicast tables, which prompts the hosts on all ports in the VLAN to reply almost simultaneously, which demands a great amount of CPU processing. It is far more desirable, from a CPU load perspective, to spread query responses out over time.

This normalization of query responses is accomplished by internally forwarding query messages to just one port at a time. This delay smooths out the load on the CPU and helps to reduce channel-change latency. This delay between ports is configured as the query delay. In addition, since the device employs IGMP Message Shaping, the router is unaware of any delay in the report (refer to the IGMP message shaping section).

Proxy Query Delay also helps to detect the path between the device and the router. If a link failure occurs that affects the path between the router and the service delivery switch, it is able to find the new path to the router even if the router does not send a query message. Instead of waiting for a query from the router, which could take minutes or even hours. The maximum time required to locate the new router path is only as long as the Proxy Query delay interval.

IGMP query engine

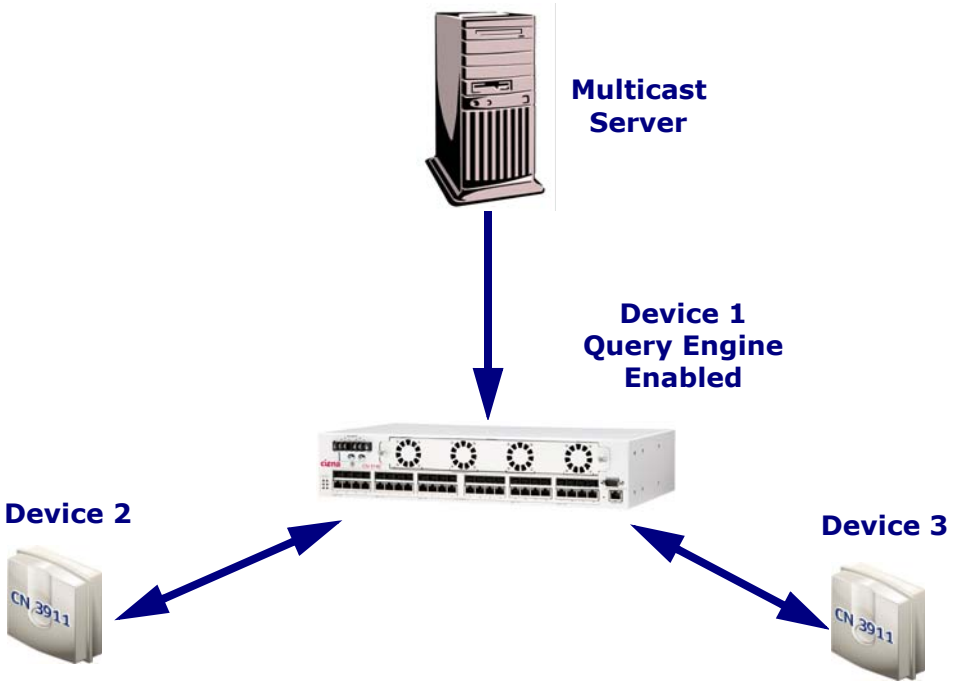
In a traditional multicast environment, all requests and queries are processed through a multicast router. Typically, these routers are overburdened with processing requirements and thus cannot respond as quickly as is needed in converged service network. To eliminate the need for a separate IGMP router and speed up multicast response times, you can configure service delivery switches to use a built-in IGMP query engine.

The query engine can be configured to generate query packets at regular intervals performing the function of a multicast router. However, only one device on the network can act as an IGMP query agent at a time and all multicast servers must interface directly with this device. Queriers negotiate with each other to determine which device should be the active querier. However, if multiple queriers are seen by the service delivery switch, it acts only on the queries from the device with the lowest IP address and ignores any other active queriers. Therefore a unique, non-zero source IP address must be set for the query engine when it is enabled.

Copyright© 2012-2013 Ciena® Corporation

Figure 14-1 is an example of an active Query Engine in a network. The multicast server interfaces directly with the “top-level” device (device 1), which is also acting as the IGMP querier for the network. The other devices forward IGMP messages to the “top-level” device only. Multicast streams flow from device 1 to device 2 and from device 1 to device 3.

Figure 14-1
Query Engine Deployment



If the query engine was enabled on device 2 instead of device 1, no membership reports would forward from device 2 to device 1, and no multicast streams would flow from device 1 to device 2.

When using the Query Engine, the device must not have a learned or configured router port for the VLAN. Additionally, you need to set the Source IP address to a value other than the default value of 0.0.0.0.

> ***Note:*** Do not enable this feature if using a traditional IGMP router on the network.

IGMP message filtering

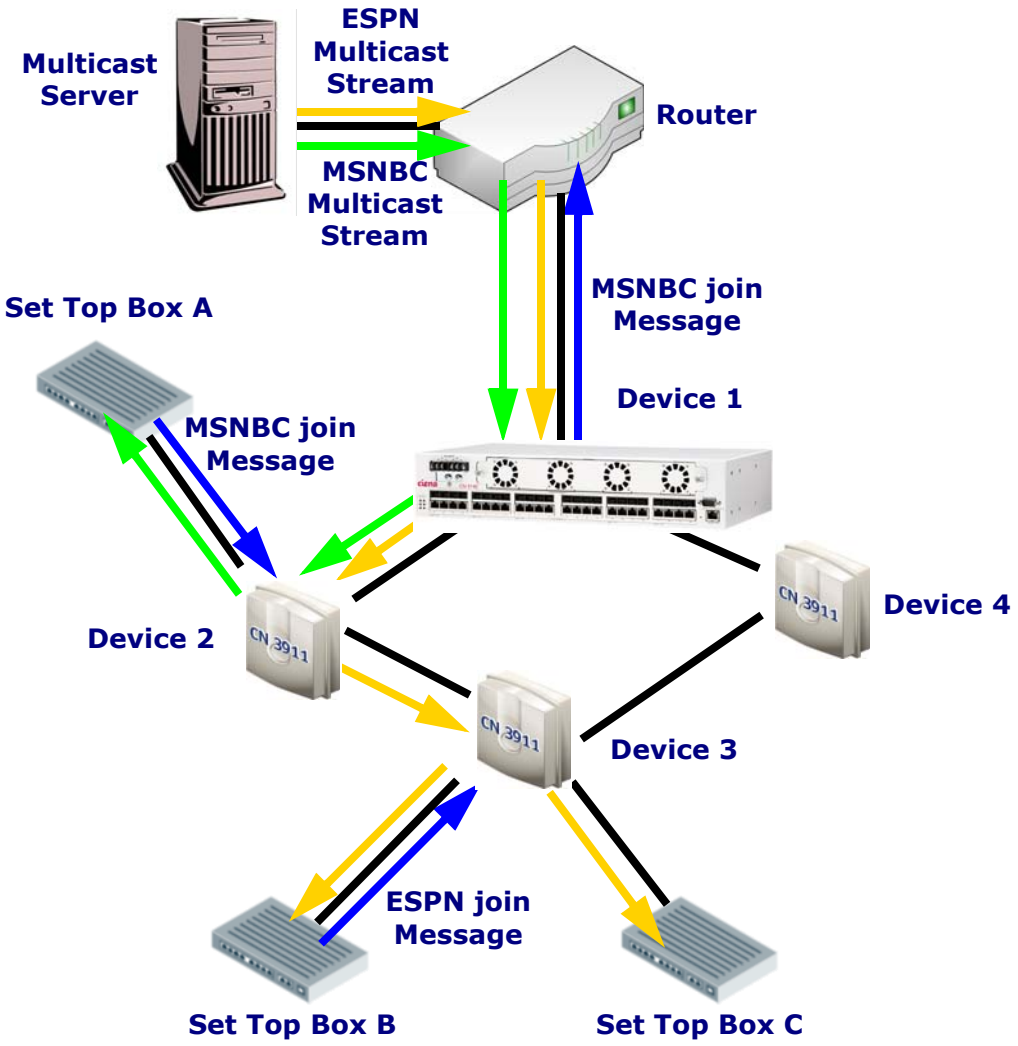
Traditional IGMP join message processing involves forwarding each IGMP join message received from a host to the server ports. Forwarding each IGMP join message to the server ports is a scaling issue for those devices higher in the network and ultimately the IGMP router. Enabling IGMP Message Filtering

Copyright© 2012-2013 Ciena® Corporation

is a scaling feature that causes the system software to intelligently forward IGMP join messages to server ports only if that IGMP join message is required to cause a multicast stream to flow to the host. IGMP join messages that are redundant are filtered.

Note: If two IGMP join messages are received on two ports, and the last 3 octets of the destination IP address are the same, the second port cannot be part of the multicast group. For example, Port 1 receives a JOIN with destination IP 224.1.2.3, and it becomes part of the multicast group. If Port 2 receives a JOIN message with destination IP 225.1.2.3, then Port 2 cannot become part of the multicast group.

Figure 14-2
IGMP Message Filtering



Copyright© 2012-2013 Ciena® Corporation

For example, if a subscriber on Set Top Box A is watching a video channel such as ESPN, when the subscriber on Set Top Box B changes the channel to ESPN (subscribes to the ESPN multicast group) Set Top Box B sends an IGMP join message to device number 3. However, device 3 does not forward the join message upstream since the multicast stream is already present. In contrast, if Set Top Box A joins the MSNBC multicast group, a join message would be sent all the way upstream back to the multicast router since the router does not have any active members in the MSNBC group.

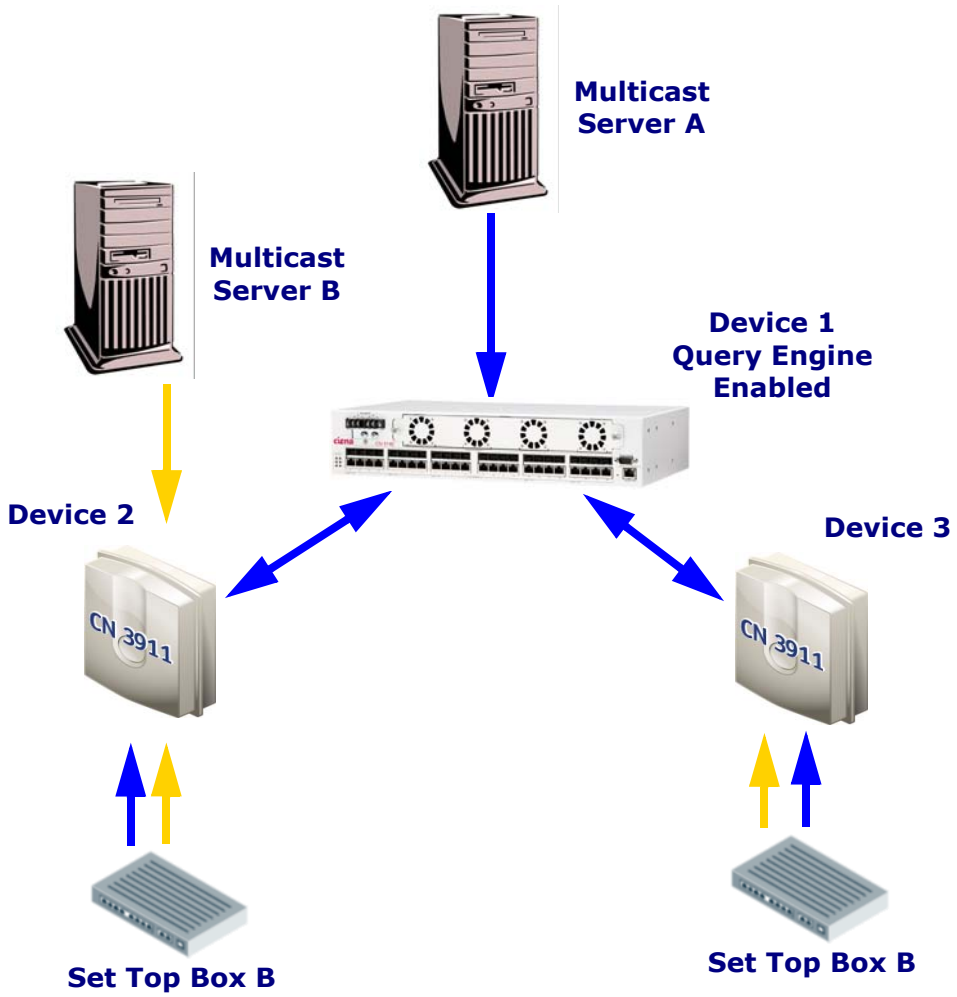
StreamCast

Typical multicast video deployments support a centralized server topology where a single video source provides all multicast content for a VLAN. While backup servers may exist, only one server delivers content.

A distributed solution allows a number of multicast servers to deliver content on a VLAN. Figure 14-3 shows such an architecture where two multicast routers each actively provide video content on the same VLAN. StreamCast is the solution that supports distributed video architecture, which allows service offerings that provide customers the ability to set up video conferences between many remote offices or other multicast streaming applications.

Copyright© 2012-2013 Ciena® Corporation

Figure 14-3
StreamCast



IGMP leave modes

Two methods for handling IGMP leaves are supported, Fast Leave and Inquisitive Leave. In Fast Leave mode, a port is removed from the multicast group immediately after a leave report is received for the port. The multicast stream no longer flows to the port.

In Inquisitive Leave mode, a timer is used to determine when ports are removed from a multicast group. Instead of immediately removing a port, when the device receives the leave report, it starts a timer and also sends out a group-specific membership query. If a membership report is received for the

Copyright© 2012-2013 Ciena® Corporation

group on that port before the timer expires, the timer is cancelled. However, if the port does not receive a membership report, the port will be removed as soon as the timer expires.

This delay, reduces channel-change latency, which can be especially useful in situations such a last channel recall. It also resolves problems introduced when a subscriber connects a hub or non-IGMP switch to their customer premises equipment (CPE). Additionally, it allows multiple set top boxes on a port to stay on their channel when another set top box changes its channel. Inquisitive leave should only be enabled on devices that have a CPE (e.g., a set top box) attached to them.

IGMP parameters

Table 14-1 lists IGMP parameters. In cases where these parameters have different names between the Device Manager and the CLI, the Device Manager name is listed in parentheses.

Table 14-1
IGMP parameters

Field	Description
IGMP Snooping (Active State)	Displays whether IGMP is active for the entry selected.
Default Router Port	The port on the device used to communicate with the router. This port sends and receives the IGMP query and join/leave messages. If a router port is not assigned, the value will be 0. Used if the actual router port is not known.
IGMP Leave Mode (Inq. Leave State)	Enables inquisitive leave. When enabled, the Active Linger Timeout is used to remove ports from the multicast group membership table. When a leave report is received by the device, the Last Member Query Interval timer is started and a group-specific query will egress the port. The timer is canceled if a membership report is received on the port. The port is removed from the group when the timer expires.
Last Member Query Interval (Last Mem. Query Int.)	Max response time inserted into group-specific queries sent in response to leave messages. Also determines the amount of time between group-specific queries. Ignored when inquisitive-leave is disabled.
L2 Packet Priority (Priority)	802.1D priority value assigned to IGMP packets for this VLAN. Default priority is 7.
Snooping State	Enables or disables IGMP snooping for the VLAN.
Robustness	Number of IGMP join or IGMP leave messages to send for each multicast group.

Copyright© 2012-2013 Ciena® Corporation

Table 14-1
IGMP parameters

Field	Description
Linger Timeout	The time, in seconds, that a filter will remain in the IGMP table after all hosts have left the group. This is done to give sufficient time for the upstream IGMP router to receive and process the leave message. The filter is removed when the Linger Timeout expires.
Active Linger Timeout	Used for inquisitive leaves. If no join message are received before the timer expires, the port is removed from the group.
Router Query Interval (Router Query Int.)	Determines the time (in seconds) to wait for the router to send a General Query message.
Minimum Query Response Time	Time, in deciseconds (1/10 of a second), that hosts have to respond to an IGMP query.
Group Address Range (Start) (Rtr. Range Start IP)	Configures the starting IP address in the multicast range for the IGMP router that corresponds to this multicast entry. Each router should be configured to use a separate range of multicast addresses to avoid conflict.
Group Address Range (End) (Rtr. Range End IP)	Configures the ending IP address in the multicast range for the IGMP router that corresponds to this multicast entry.
Proxy Query Source Address (Source IP Address)	When the query engine is enabled, this address is used as the source for all query packets. This IP address must be unique for the VLAN.
IGMP Query Engine (Query Engine State)	Enables/Disables the IGMP query engine.
Proxy Query Interval (Query Interval)	Determines the interval, in seconds when the query engine will send IGMP general queries.
Proxy Query Response Interval	Determines the rate at which devices respond to IGMP queries from the router.
Proxy Query Delay	Query messages are sent to one port at a time. This setting determines the delay between forwarding messages to each port.

Unknown multicast filtering

Various filters are used to limit multicast traffic. Normally when a Layer-2 switch receives a multicast packet with an unknown destination address, it floods the packet to all ports in the VLAN. Unknown multicast filtering (UMF) is used to filter (drop) these unknown packets. A destination address is considered unknown if it is not found in the list of active multicast groups.

Copyright© 2012-2013 Ciena® Corporation

UMF is useful when a multicast server is connected directly to the Layer-2 network, bypassing the Layer-3 router. In this case, the multicast server continuously floods the device with multicast traffic. With UMF enabled, the switch will simply drop all multicast streams that have not been joined until a a join message is received. UMF is enabled by setting this feature to Filter (drop). UMF should be enabled before reaching host devices otherwise, they will be flooded with multiple multicast streams.

Configuring multicast

The following examples show various IGMP configurations in ring topologies.The default settings for the IGMP timers should work well for most applications.

In addition, there are two important configuration settings on the multicast router that can affect IGMP performance.

- **Query-Response-Interval** -determines how fast the device must respond to a general query. A short interval results in a burst of responses within a very short time, causing the CPU load to spike. A larger value helps avoid this peak thereby easing CPU load. Where possible, this value should be set as close to (but not exceeding) the query-interval.
- **Query-Interval** -on the router, this value determines how often the router sends a general-query message. Each time the device receives a general query, it must use system resources to respond to the router. If the interval is too short it could cause the device to become sluggish. The minimum value should be about one minute for every 100 multicast groups on the network. The value chosen for query-interval on the router should be less than or equal to the value chosen for router-query-interval on the device.

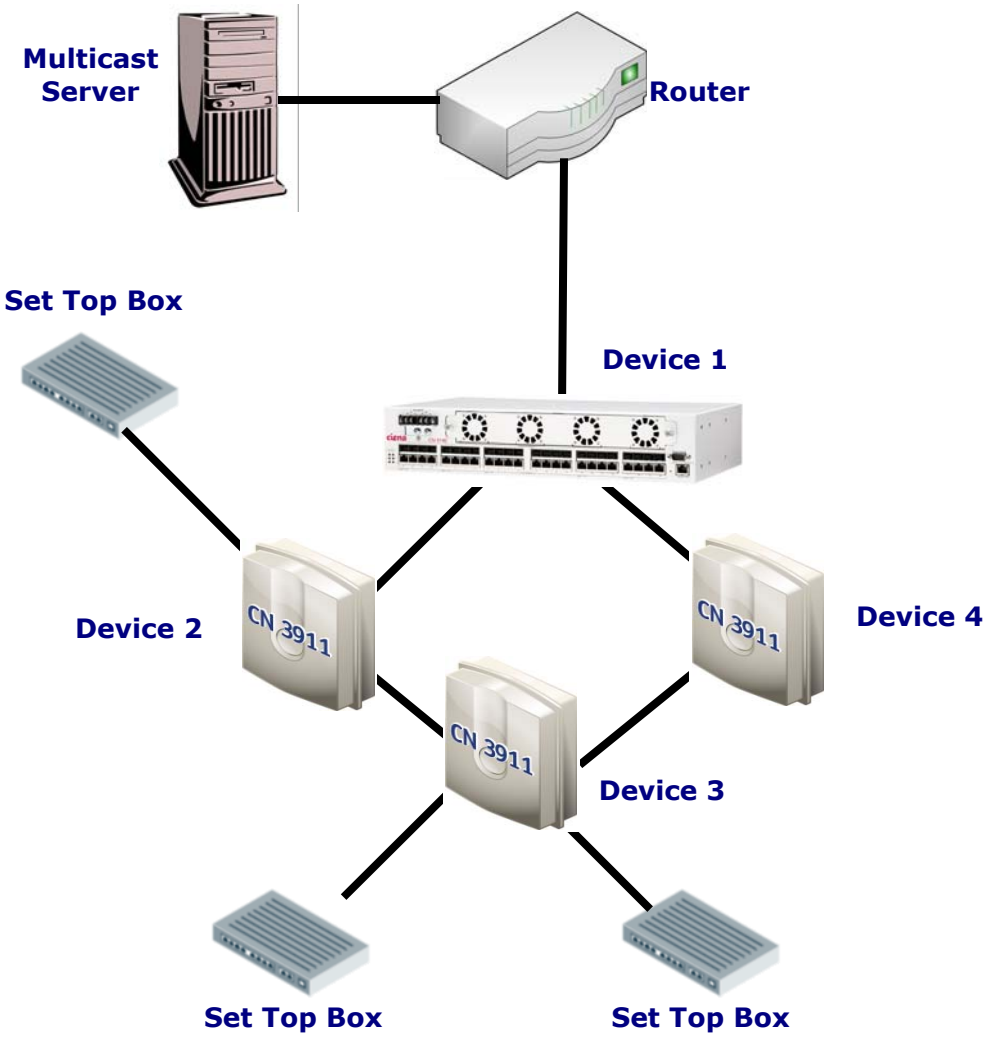
Copyright© 2012-2013 Ciena® Corporation

Procedure 14-1
Configuring a multicast router topology

This example demonstrates how to configure devices in a network where a multicast server is delivering services over the network via a multicast router. This router is also the querier and controls the delivery of multicast streams.

In this example, multicast services are delivered on VLAN 300.

Figure 14-4
Multicast Services with a Multicast Router



The default value for the query-ip-source-address is 0.0.0.0. This value is used as the source address for all query messages as well as for proxy-reply messages sent to the router. The default value should work in most cases.

Copyright© 2012-2013 Ciena® Corporation

However, if this address must be changed because of a router on the network, make sure to use the correct address for the router interface you are using, otherwise, the router may discard all proxy-reply messages from the device.

Step	Action
1	Create a multicast entry for the VLAN on all four devices, which enables multicast services for the specified VLAN: `multicast-services create vlan <VlanList>`
2	Set the query timers to the same values as on the multicast router. `multicast-services igmp-snooping set`
3	Enable IGMP snooping on all four devices. `multicast-services igmp-snooping enable`
4	To enable UMF on the specified VLAN for the devices that have subscriber equipment connected to them (in our example, device 2 and 3), enter the following command: `multicast-services umf drop` In this scenario, UMF blocks multicast streams originating from subscriber equipment or other locations in the network. It also ensures the subscriber devices receive only the multicast streams they have joined.
5	To verify the IGMP settings, enter the following command: `multicast-services show configuration`

—end—

Example

The following example configures device 2 as seen in Figure 14-4.

```
multicast-services create vlan 300
multicast-services igmp-snooping set query-response-
interval 100
multicast-services igmp-snooping enable
multicast-services umf drop vlan 300
multicast-services wkm-forwarding add vlan 300 wkm-group
internetwork-control
multicast-services show configuration vlan 300
```

```
+-----------------------------------------------------+-------------------------+
|    MULTICAST CONFIGURATION VLAN 300 VLAN#300                                  |
+-----------------------------------------------------+-------------------------+
| Global Multicast Services                           | enable                  |
| Multicast Services                                  | enable                  |
| UMF                                                 | drop                    |
| WKM Forwarding                                      | disabled                |
+-----------------------------------------------------+-------------------------+
|    IGMP CONFIGURATION (general)                                               |
+-----------------------------------------------------+-------------------------+
| IGMP Snooping                                       | enable                  |
```


Copyright© 2012-2013 Ciena® Corporation

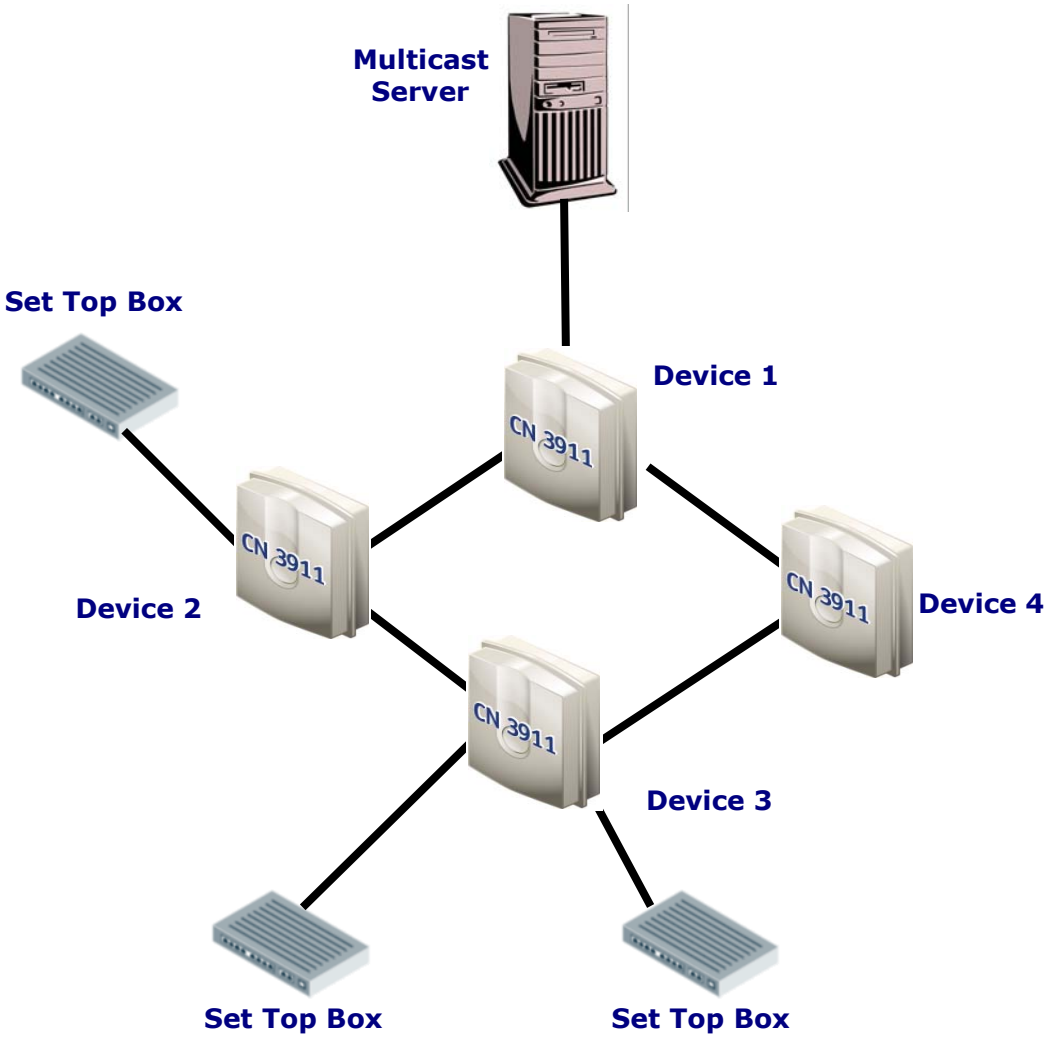
IGMP Server Topology	centralized
IGMP Query Engine	off
IGMP Leave Mode	fast
L2 Packet Priority	7
Robustness	1
IGMP CONFIGURATION (proxy query)	
Proxy Query Interval (s)	125
Proxy Query Response Interval (ds)	100
Proxy Query Delay (ds)	10
Proxy Query Source Address	0.0.0.0
Last Member Query Interval (ds)	10
IGMP CONFIGURATION (router)	
Router Query Interval (s)	250
Linger Timeout (s)	120
Active Linger Timeout (s)	30
Minimum Query Response Interval (ds)	50
Group Address Range (start)	0.0.0.0
Group Address Range (end)	0.0.0.0
Default Router Port	0

Copyright© 2012-2013 Ciena® Corporation

Procedure 14-2
Configuring a multicast server topology

This example shows how to configure devices in a network that does not employ a multicast router. The multicast server floods a set of multicast streams on the network. The device connected to the multicast server is configured to contain the multicast streams and to act as the querier (Query Engine enabled) for downstream devices. For this example, multicast services are delivered on VLAN 300. In the topology represented in Figure 14-5, UMF is required on device 1 to contain the streams being flooded by the multicast server. On all other devices, UMF is optional.

Figure 14-5
Delivering Multicast Services using the Query Engine



Copyright© 2012-2013 Ciena® Corporation

Step	Action
1	Create a multicast entry for the VLAN on all four devices. `multicast-services create vlan <VlanList>` This enables multicast services for the specified VLAN.
2	Configure the Query Engine on the same device. `multicast-services igmp-snooping set` When using this command, the source IP address must be set before the Query Engine is enabled. In addition, the router port must be set to 0 (unlearned).
3	Enable IGMP snooping on all four devices. `multicast-services igmp-snooping enable`
4	Enable UMF on device 1. (Optional) Enable UMF on the specified VLAN for the devices that have subscriber equipment connected to them (in our example, device 2 and 3). `multicast-services umf drop vlan <VlanList>`
5	To verify the IGMP settings, enter the following command: `multicast-services show configuration`

—end—

Example

The following example configures device 1 as shown in Figure 14-5.

```
multicast-services create vlan 300
multicast-services igmp-snooping set query-ip-source-
  addr 10.10.10.10 query-engine on
multicast-services igmp-snooping enable
multicast-services umf drop vlan 300
multicast-services wkm-forwarding add vlan 300 wkm-group
internetwork-control
multicast-services show configuration vlan 300
```

```
+-----------------------------------------------------+-------------------------+
|    MULTICAST CONFIGURATION VLAN 300 VLAN#300                                  |
+-----------------------------------------------------+-------------------------+
| Global Multicast Services                           | enable                  |
| Multicast Services                                  | enable                  |
| UMF                                                 | drop                    |
| WKM Forwarding                                      | disabled                |
+-----------------------------------------------------+-------------------------+
|    IGMP CONFIGURAITON (general)                                               |
+-----------------------------------------------------+-------------------------+
| IGMP Snooping                                       | enable                  |
| IGMP Query Engine                                   | on                      |
| IGMP Leave Mode                                     | fast                    |
| L2 Packet Priority                                  | 7                       |
| Robustness                                          | 1                       |
+-----------------------------------------------------+-------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

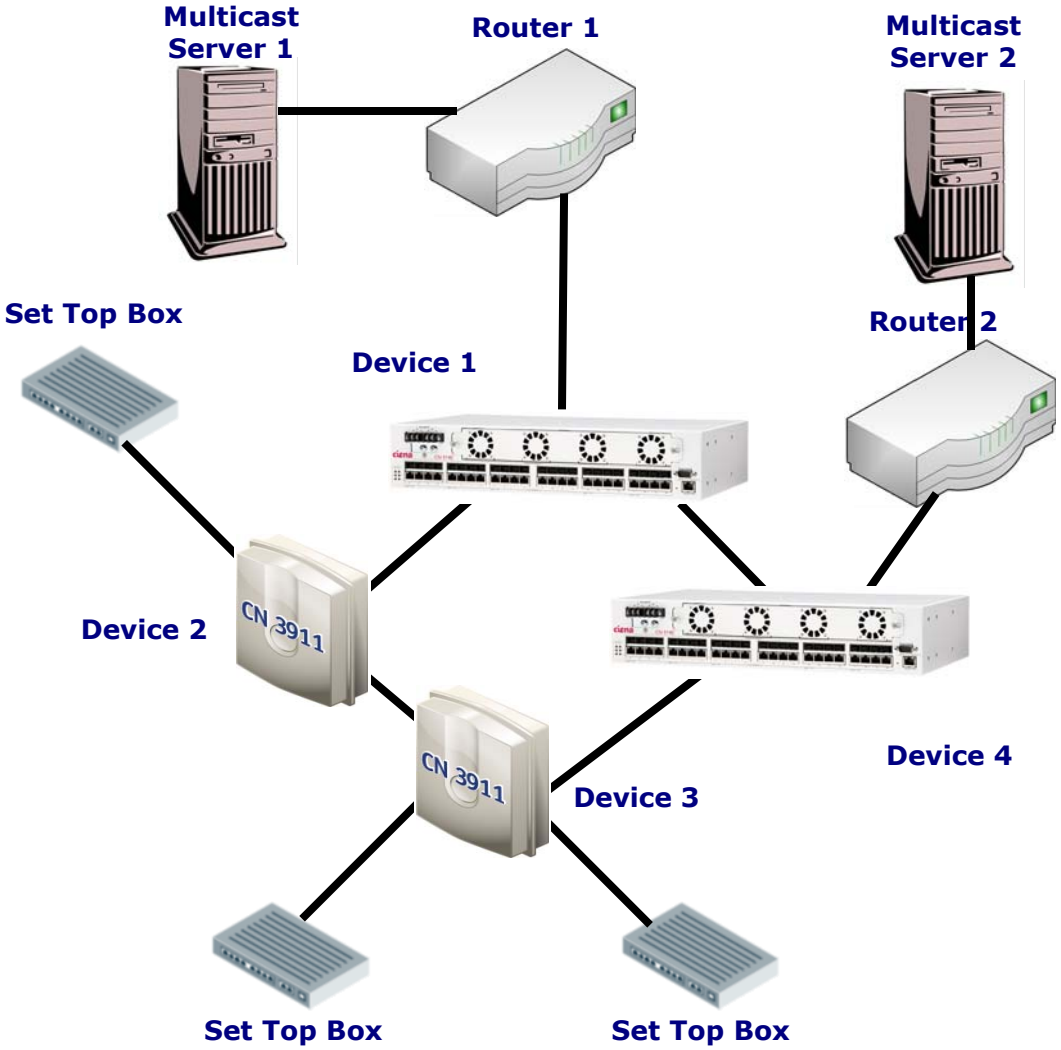
IGMP CONFIGURAITON (proxy query)	
Proxy Query Interval (s)	125
Proxy Query Response Interval (ds)	100
Proxy Query Delay (ds)	10
Proxy Query Source Address	10.10.10.10
Last Member Query Interval (ds)	10
IGMP CONFIGURAITON (router)	
Router Query Interval (s)	250
Linger Timeout (s)	120
Active Linger Timeout (s)	30
Minimum Query Response Interval (ds)	50
Group Address Range (start)	0.0.0.0
Group Address Range (end)	0.0.0.0
Default Router Port	0

Copyright© 2012-2013 Ciena® Corporation

Procedure 14-3
Configuring multicast servers with redundant routers

This example demonstrates how to configure devices in a network that has two (or more) redundant multicast servers. Both multicast servers are connected to the network via a multicast router. In this scenario, only one router with its associated multicast server, is active at a time. The second router and server in a standby mode. The active server is determined by identifying the router that has the lowest IP address. This router will handle general queries and “authorize” the delivery of various streams through the network.

Figure 14-6
Delivering Multicast Services Using Redundant Multicast Routers



In response to a topology change, for example, if the link that has the router port for the device goes down, device will advertise a router address of 0.0.0.0 to indicate that it has lost the router. For example, assuming that router 1 has a lower IP address than Router 2, if the link between device 1 and Router 1 goes down, devices 1, 2, 3 and 4 will send out general queries advertising a router address of 0.0.0.0. Device 4 will become the interface for multicast services and all joins/leaves will then be sent to it. However, when Router 1 is restored, it will again become the primary router and Router 2 will again go into standby mode.

This allows the network to continue delivery of multicast services while ensuring the least delay in service possible. This also allows order to be maintained between the two multicast routers, so that there is no competition or confusion.

In this example, multicast services are assumed to be delivered on VLAN 300. For the network represented in Figure 14-6, UMF is required on device 1 and device 2. On all other devices, it is optional.

The default value for the query-ip-source-address is 0.0.0.0. This value is used as the source address for all query messages as well as for proxy-reply messages sent to the router. The default value should work in most cases. However, if this address must be changed because of a router on the network, make sure to use the correct address for the router interface you are using, otherwise, the router may discard all proxy-reply messages from the device.

Step	Action
1	Create a multicast entry for the VLAN on all four devices. `multicast-services create vlan <VlanList>` This enables multicast services for the specified VLAN.
2	Enable UMF on device 1 and device 4. (Optional) Enable UMF on the specified VLAN for the devices that have subscriber equipment connected to them (in our example, device 2 and 3). `multicast-services umf drop`
3	Configure WMF to allow desirable streams through on device 1 and device 4. `multicast-services wkm-forwarding add <VLAN-object-list> wkm-group <'adhoc' | 'internetwork-control' | 'local-control' | 'sdp-sap' | 'st-multicast'>`
4	Set the query timers to the same values as on the multicast router. `multicast-services igmp-snooping set`
5	Enable IGMP snooping on all four devices. `multicast-services igmp-snooping enable`

Copyright© 2012-2013 Ciena® Corporation

6 If UMF is enabled, configure WMF to allow desirable streams through.

```
multicast-services wkm flood vlan <vlan-name> <'adhoc' |
'internetwork-control' | 'local-control' | 'sdp-sap' |
'st-multicast'>
```

7 (Optional) Verify the IGMP settings.

```
multicast-services show configuration
```

—end—

Example

The following example configures device 1 as shown in Figure 14-6. Note that device 1 and device 4 have the same configuration.

```
multicast-services create vlan 300
multicast-services umf drop vlan 300
multicast-services wkm-forwarding add vlan 300 wkm-group
internetwork-control
multicast-services igmp-snooping set query-response-
interval 100
multicast-services igmp-snooping enable
multicast-services show configuration vlan 300
```

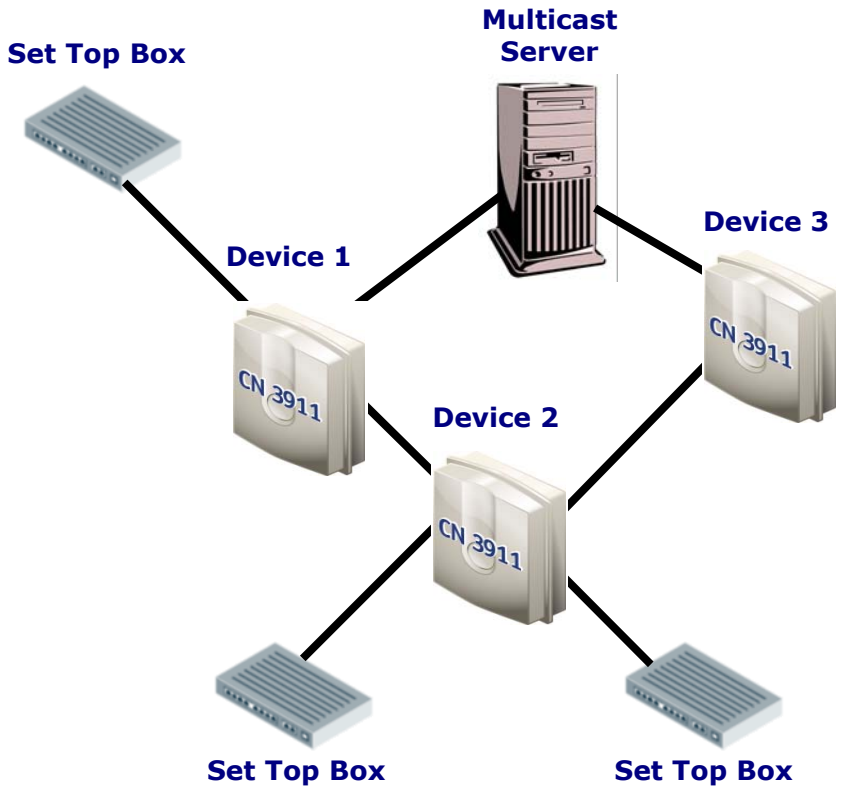
```
+-----------------------------------------------------+-------------------------+
|    MULTICAST CONFIGURATION VLAN 300 VLAN#300                                  |
+-----------------------------------------------------+-------------------------+
| Global Multicast Services                           | enable                  |
| Multicast Services                                  | enable                  |
| UMF                                                 | drop                    |
| WKM Forwarding                                      | disabled                |
+-----------------------------------------------------+-------------------------+
|    IGMP CONFIGURATION (general)                                               |
+-----------------------------------------------------+-------------------------+
| IGMP Snooping                                       | enable                  |
| IGMP Server Topology                                | centralized             |
| IGMP Query Engine                                   | off                     |
| IGMP Leave Mode                                     | fast                    |
| L2 Packet Priority                                  | 7                       |
| Robustness                                          | 1                       |
+-----------------------------------------------------+-------------------------+
|    IGMP CONFIGURATION (proxy query)                                           |
+-----------------------------------------------------+-------------------------+
| Proxy Query Interval (s)                            | 125                     |
| Proxy Query Response Interval (ds)                  | 100                     |
| Proxy Query Delay (ds)                              | 10                      |
| Proxy Query Source Address                          | 0.0.0.0                 |
| Last Member Query Interval (ds)                     | 10                      |
+-----------------------------------------------------+-------------------------+
|    IGMP CONFIGURATION (router)                                                |
+-----------------------------------------------------+-------------------------+
| Router Query Interval (s)                           | 250                     |
| Linger Timeout (s)                                  | 120                     |
| Active Linger Timeout (s)                           | 30                      |
| Minimum Query Response Interval (ds)                | 50                      |
| Group Address Range (start)                         | 0.0.0.0                 |
| Group Address Range (end)                           | 0.0.0.0                 |
| Default Router Port                                 | 0                       |
+-----------------------------------------------------+-------------------------+
```


Copyright© 2012-2013 Ciena® Corporation

Procedure 14-4
Configuring redudant query engines

This example demonstrates how to configure devices in a network that has redundant links between the multicast server and the rest of the network.

Figure 14-7
IGMP Snooping



In this scenario, the Query Engine must be enabled on device 1 and device 3, but only one Query Engine will be active at a time. The second Query Engine will effectively be in standby mode. The active Query Engine is determined by identifying the engine with the lowest source IP address. This Query Engine then handles general queries and “authorizes” the delivery of various streams through the network.

In response to a topology change, for example, if the link that has the router port for a device goes down, the device will advertise a router address of 0.0.0.0 to indicate that it has lost the router.

Copyright© 2012-2013 Ciena® Corporation

For example, assuming that device 1 has a lower source IP address than device 3, if the link between device 1 and the Multicast Server goes down, devices 1, 2, and 3 will send out general queries advertising a router address of 0.0.0.0. Device 3 will then become the interface for multicast services and all joins/leaves will then be sent to it. However, when the link between device 1 and the Multicast Server is restored, it will again become the primary Query Engine, and device 3 will return to standby mode.

This allows the network to continue delivery of multicast services while ensuring the least delay in service possible. This also allows order to be maintained between the two Query Engines, so that there is no competition or confusion.

In this example, multicast services are assumed to be delivered on VLAN 300. The network should also be using RSTP. For the network represented in Figure 14-7, UMF is required on device 1 and device 3. On all other devices, it is optional.

Step	Action
1	Create a multicast entry for the VLAN on devices 1, 2, and 3. `multicast-services create vlan <VlanList>` This enables multicast services for the specified VLAN.
2	Configure WMF to allow desirable streams through on devices 1 and 3. `multicast-services umf drop`
3	Configure the Query Engine on devices 1 and 3. `multicast-services igmp-snooping set [query-ip-source-addr <IpAddress>] [query-engine <on|off>]` When using this command, the source IP address must come before the Query Engine is enabled, in order to function properly. In addition, the router port must be set to 0 (unlearned). Make sure that the source IP address for the two devices is different and that the device that should become the primary Query Engine has a lower IP address.
4	Enable IGMP snooping on devices 1, 2, and 3. `multicast-services igmp-snooping set`
5	If UMF is enabled, configure WMF to allow desirable streams through. `multicast-services igmp-snooping enable`
6	Enable UMF on devices 1 and 3. (Optional) Enable UMF on devices that have subscriber equipment connected to them (in our example, device 2). `multicast-services umf drop`

Copyright© 2012-2013 Ciena® Corporation

7 (Optional) Verify the IGMP settings.

```
multicast-services show configuration
```

—end—

Example

The following example configures device 1 as shown in Figure 14-5. Note that with the exception of the Query Engine source IP address, the configuration for device 1 and device 3 is the same.

```
multicast-services create vlan 300
multicast-services igmp-snooping set query-ip-source-addr
10.10.10.10 query-engine on
multicast-services igmp-snooping enable
multicast-services umf drop vlan 300
multicast-services wkm-forwarding add vlan 300 wkm-group
internetwork-control
multicast-services show configuration vlan 300
```

```
+-----------------------------------------------------+-------------------------+
|    MULTICAST CONFIGURAITON VLAN 300 VLAN#300                                  |
+-----------------------------------------------------+-------------------------+
| Global Multicast Services                           | enable                  |
| Multicast Services                                  | enable                  |
| UMF                                                 | drop                    |
| WKM Forwarding                                      | disabled                |
+-----------------------------------------------------+-------------------------+
|    IGMP CONFIGURAITON (general)                                               |
+-----------------------------------------------------+-------------------------+
| IGMP Snooping                                       | enable                  |
| IGMP Query Engine                                   | on                      |
| IGMP Leave Mode                                     | fast                    |
| L2 Packet Priority                                  | 7                       |
| Robustness                                          | 1                       |
+-----------------------------------------------------+-------------------------+
|    IGMP CONFIGURAITON (proxy query)                                           |
+-----------------------------------------------------+-------------------------+
| Proxy Query Interval (s)                            | 125                     |
| Proxy Query Response Interval (ds)                  | 100                     |
| Proxy Query Delay (ds)                              | 10                      |
| Proxy Query Source Address                          | 10.10.10.10             |
| Last Member Query Interval (ds)                     | 10                      |
+-----------------------------------------------------+-------------------------+
|    IGMP CONFIGURAITON (router)                                                |
+-----------------------------------------------------+-------------------------+
| Router Query Interval (s)                           | 250                     |
| Linger Timeout (s)                                  | 120                     |
| Active Linger Timeout (s)                           | 30                      |
| Minimum Query Response Interval (ds)                | 50                      |
| Group Address Range (start)                         | 0.0.0.0                 |
| Group Address Range (end)                           | 0.0.0.0                 |
| Default Router Port                                 | 0                       |
+-----------------------------------------------------+-------------------------+
```

Copyright© 2012-2013 Ciena® Corporation

Appendix A: Error codes

This appendix lists error codes.

Traffic profiling error codes

Table A-1 describes the overlap error codes that can occur in advanced mode when provisioning multiple traffic profiles on the same port. For example, if VLAN^{A} is provisioned with the VLAN attribute value A, it will overlap with a traffic profile on the same port that has the same attribute value for A.

The following is an example of configuring an overlap condition, and the resulting error message:

```
> traffic-profiling set port 1 mode advanced
> traffic-profiling standard-profile create port 1 profile 1 name t1 cir 1024
pir 1024 cbs 32 ebs 0 dot1dpri 1
> traffic-profiling standard-profile create port 1 profile 2 name t2 cir 1024
pir 2048 cbs 32 ebs 0 dot1dpri 1

ERROR: Creating standard-profile 2 overlaps profile 1 on port 1 -- classifier
overlaps existing profile overlap_code #28
```

Table A-1
Traffic profiling error codes

Overlap code	Classifier combination	Cause of overlap
1	Untagged	Profile with either IP or dscp classifier
2	Vlan^{A}	Profile with VS + Vlan classifiers and the Vlan ^{A} is a member of the VS. Profile with VS +Vlan + any COS (IP, dscp, dot1d) classifiers and the Vlan ^{A} is a member of the VS.
3	Vlan^{A}	Profile with VS classifier and Vlan^{A} is a member of the VS. Profile with VS + any COS (IP,dscp,dot1d) classifiers and Vlan^{A} is a member of the VS.
4	Vlan^{A}	Profile with Vlan^{A} classifier. Profile with Vlan^{A} + any COS (IP, dscp, dot1d) classifiers.

Copyright© 2012-2013 Ciena® Corporation

Table A-1
Traffic profiling error codes

5	Vlan[A]	Profile with any COS (IP, dscp, dot1d) classifier.
6	Vlan[A] + dot1d[B]	Profile with VS + dot1d[B] classifiers and Vlan[A] is a member of the VS and dot1d value matches.
7	Vlan[A] + dot1d[B]	Profile with VS+Vlan[A]+dot1d[B] classifiers and Vlan[A] is a member of the VS and the dot1d value matches.
8	Vlan[A] + dot1d[B]	Profile with VS classifier and Vlan[A] is a member of the VS. Profile with VS + IP or dscp classifiers and Vlan[A] is a member of the VS. Profile with VS + Vlan classifiers and Vlan[A] is a member of the VS. Profile with VS + Vlan +IP or dscp classifiers and Vlan[A] is a member of the VS.
9	Vlan[A] + dot1d[B]	Profile with Vlan[A] +dot1d[B] classifiers and dot1d value matches.
10	Vlan[A] + dot1d[B]	Profile with Vlan[A] classifier. Profile with Vlan[A] + IP or dscp classifiers.
11	Vlan[A] + dot1d[B]	Profile with IP or dscp classifiers.
12	Vlan[A] + dot1d[B]	Profile with dot1d[B] classifier and dot1d value matches.
13	Vlan[A] + IP[B]	Profile with VS + IP[B] classifiers and Vlan[A] is a member of the VS and IP value matches.
14	Vlan[A] + IP[B]	Profile with VS+Vlan[A]+IP[B] classifiers and Vlan[A] is a member of the VS and the IP value matches.
15	Vlan[A] + IP[B]	Profile with VS classifier and Vlan[A] is a member of the VS. Profile with VS + dot1d or dscp classifiers and Vlan[A] is a member of the VS. Profile with VS + Vlan classifiers and Vlan[A] is a member of the VS. Profile with VS + Vlan +dot1d or dscp classifiers and Vlan[A] is a member of the VS.
16	Vlan[A] + IP[B]	Profile with Vlan[A] +IP[B] classifiers and IP value matches.
17	Vlan[A] + IP[B]	Profile with Vlan[A] classifier. Profile with Vlan[A] + dot1d or dscp classifiers.
18	Vlan[A] + IP[B]	Profile with dot1d or dscp classifiers.
19	Vlan[A] + IP[B]	Profile with IP[B] classifier and IP value matches.

Copyright© 2012-2013 Ciena® Corporation

Table A-1
Traffic profiling error codes

20	Vlan[A] + dscp[B]	Profile with VS + dscp[B] classifiers and Vlan[A] is a member of the VS and dscp value matches.
21	Vlan[A] + dscp[B]	Profile with VS+Vlan[A]+dscp[B] classifiers and Vlan[A] is a member of the VS and the dscp value matches.
22	Vlan[A] + dscp[B]	Profile with VS classifier and Vlan[A] is a member of the VS. Profile with VS + IP or dot1d classifiers and Vlan[A] is a member of the VS. Profile with VS + Vlan classifiers and Vlan[A] is a member of the VS. Profile with VS + Vlan +IP or dot1d classifiers and Vlan[A] is a member of the VS.
23	Vlan[A] + dscp[B]	Profile with Vlan[A] +dscp[B] classifiers and dscp value matches.
24	Vlan[A] + dscp[B]	Profile with Vlan[A] classifier. Profile with Vlan[A] + IP or dot1d classifiers.
25	Vlan[A] + dscp[B]	Profile with IP or dot1d classifiers.
26	Vlan[A] + dscp[B]	Profile with dscp[B] classifier and dscp value matches.
27	dot1d[A]	Profile with VS classifier. Profile with VS + IP or dscp classifiers. Profile with VS + Vlan classifiers. Profile with VS + Vlan + IP or dscp classifiers. Profile with Vlan classifier. Profile with Vlan + IP or dscp classifiers. Profile with IP or dscp classifier.
28	dot1d[A]	Profile with VS + dot1d[A] classifiers. Profile with VS + Vlan +dot1d[A] classifiers. Profile with Vlan +dot1d[A] classifiers. Profile with dot1d[A] classifier. Where the dot1d value matches.

Copyright© 2012-2013 Ciena® Corporation

Table A-1
Traffic profiling error codes

29	IP[A]	Profile with VS classifier. Profile with VS + dot1d or dscp classifiers. Profile with VS + Vlan classifiers. Profile with VS + Vlan + dot1d or dscp classifiers. Profile with Vlan classifier. Profile with Vlan + dot1d or dscp classifiers. Profile with dot1d or dscp classifier.
30	IP[A]	Profile with VS + IP[A] classifiers. Profile with VS + Vlan +IP[A] classifiers. Profile with Vlan +IP[A] classifiers. Profile with IP[A] classifier. Where the IP value matches.
31	dscp[A]	Profile with VS classifier. Profile with VS + IP or dot1d classifiers. Profile with VS + Vlan classifiers. Profile with VS + Vlan + IP or dot1d classifiers. Profile with Vlan classifier. Profile with Vlan + IP or dot1d classifiers. Profile with IP or dot1d classifier.
32	dscp[A]	Profile with VS + dscp[A] classifiers. Profile with VS + Vlan +dscp[A] classifiers. Profile with Vlan +dscp[A] classifiers. Profile with dscp[A] classifier. Where the dscp value matches.
33	VS[A]	Profile with VS[A] classifier. Profile with VS[A] + any COS(dot1d,IP,dscp) classifiers. Profile with VS[A] +Vlan classifiers. Profile with VS[A] + Vlan + any COS(dot1d,IP,dscp) classifiers.
34	VS[A]	Profile with Vlan classifier and Vlan is a member of VS[A]. Profile with Vlan + any COS(dot1d, IP, dscp) classifiers and Vlan is a member of VS[A].
35	VS[A]	Profile with any COS(dot1d,IP,dscp) classifier.

Copyright© 2012-2013 Ciena® Corporation

Table A-1
Traffic profiling error codes

36	VS^A + $dot1d^B$	Profile with VS^A + $dot1d^B$ classifiers and dot1d value matches. Profile with VS^A +Vlan + $dot1d^B$ classifiers and dot1d value matches.
37	VS^A + $dot1d^B$	Profile with VS^A classifier. Profile with VS^A +IP or dscp classifiers. Profile with VS^A +Vlan classifiers. Profile with VS^A + Vlan + IP or dscp classifiers.
38	VS^A + $dot1d^B$	Profile with Vlan +$dot1d^B$ classifiers and Vlan is a member of VS^A.
39	VS^A + $dot1d^B$	Profile with Vlan and Vlan is a member of VS^A. Profile with Vlan + IP or dscp classifiers and Vlan is a member of VS^A.
40	VS^A + $dot1d^B$	Profile with $dot1d^B$ classifier.
41	VS^A + $dot1d^B$	Profile with IP or dscp classifier.
42	VS^A + IP^B	Profile with VS^A + IP^B classifiers and IP value matches. Profile with VS^A +Vlan + IP^B classifiers and IP value matches.
43	VS^A + IP^B	Profile with VS^A classifier. Profile with VS^A +dot1d or dscp classifiers. Profile with VS^A +Vlan classifiers. Profile with VS^A + Vlan + dot1d or dscp classifiers.
44	VS^A + IP^B	Profile with Vlan +IP^B classifiers and Vlan is a member of VS^A.
45	VS^A + IP^B	Profile with Vlan and Vlan is a member of VS^A. Profile with Vlan + dot1d or dscp classifiers and Vlan is a member of VS^A.
46	VS^A + IP^B	Profile with IP^B classifier.
47	VS^A + IP^B	Profile with dot1d or dscp classifier.
48	VS^A + $dscp^B$	Profile with VS^A + $dscp^B$ classifiers and dscp value matches. Profile with VS^A +Vlan + $dscp^B$ classifiers and dscp value matches.
49	VS^A + $dscp^B$	Profile with VS^A classifier. Profile with VS^A +IP or dot1d classifiers. Profile with VS^A +Vlan classifiers. Profile with VS^A + Vlan + IP or dot1d classifiers.
50	VS^A + $dscp^B$	Profile with Vlan +$dscp^B$ classifiers and Vlan is a member of VS^A.

Copyright© 2012-2013 Ciena® Corporation

Table A-1
Traffic profiling error codes

51	VS^A + $dscp^B$	Profile with Vlan and Vlan is a member of VS^A. Profile with Vlan + IP or dot1d classifiers and Vlan is a member of VS^A.
52	VS^A + $dscp^B$	Profile with $dscp^B$ classifier.
53	VS^A + $dscp^B$	Profile with IP or dot1d classifier.
54	VS^A + $Vlan^B$	Profile with VS^A classifier. Profile with VS^A + any COS(dot1d, IP, dscp) classifiers.
55	VS^A + $Vlan^B$	Profile with VS^A + $Vlan^B$ classifiers. Profile with VS^A + $Vlan^B$ +any COS(dot1d,IP,dscp) classifiers.
56	VS^A + $Vlan^B$	Profile with $Vlan^B$ classifier and $Vlan^B$ is a member of VS^A. Profile with $Vlan^B$ + any COS (dot1d,IP,dscp) classifiers and $Vlan^B$ is a member of VS^A.
57	VS^A + $Vlan^B$	Profile with any COS(dot1d, IP, dscp).
58	VS^A + $Vlan^B$ + $dot1d^C$	Profile with VS^A + $dot1d^C$ classifiers and dot1d value matches.
59	VS^A + $Vlan^B$ + $dot1d^C$	Profile with VS^A classifier. Profile with VS^A +IP or dscp classifiers.
60	VS^A + $Vlan^B$ + $dot1d^C$	Profile with VS^A + $Vlan^B$ + $dot1d^C$ classifiers and dot1d value matches.
61	VS^A + $Vlan^B$ + $dot1d^C$	Profile with VS^A + $Vlan^B$ classifiers. Profile with VS^A + $Vlan^B$ + IP or dscp classifiers.
62	VS^A + $Vlan^B$ + $dot1d^C$	Profile with $Vlan^B$ +$dot1d^C$ classifiers and $Vlan^B$ is a member of VS^A and dot1d value matches.
63	VS^A + $Vlan^B$ + $dot1d^C$	Profile with $Vlan^B$ classifier and $Vlan^B$ is a member of VS^A. Profile with $Vlan^B$ + IP or dscp classifiers and $Vlan^B$ is a member of VS^A.
64	VS^A + $Vlan^B$ + $dot1d^C$	Profile with IP or dscp classifier.
65	VS^A + $Vlan^B$ + $dot1d^C$	Profile with $dot1d^C$ and dot1d value matches.
66	VS^A + $Vlan^B$ + IP^C	Profile with VS^A +IP^C classifiers and IP value matches.
67	VS^A + $Vlan^B$ + IP^C	Profile with VS^A classifier. Profile with VS^A +dot1d or dscp classifiers.
68	VS^A + $Vlan^B$ + IP^C	Profile with VS^A + $Vlan^B$ + IP^C classifiers and IP value matches.
69	VS^A + $Vlan^B$ + IP^C	Profile with VS^A + $Vlan^B$ classifiers. Profile with VS^A + $Vlan^B$ + dot1d or dscp classifiers.

Copyright© 2012-2013 Ciena® Corporation

Table A-1
Traffic profiling error codes

70	VS^A + $Vlan^B$ + IP^C	Profile with $Vlan^B$ +IP^C classifiers and $Vlan^B$ is a member of VS^A and IP value matches.
71	VS^A + $Vlan^B$ + IP^C	Profile with $Vlan^B$ classifier and $Vlan^B$ is a member of VS^A. Profile with $Vlan^B$ + dot1d or dscp classifiers and $Vlan^B$ is a member of VS^A.
72	VS^A + $Vlan^B$ + IP^C	Profile with dot1d or dscp classifier.
73	VS^A + $Vlan^B$ + IP^C	Profile with IP^C and IP value matches.
74	VS^A + $Vlan^B$ + $dscp^C$	Profile with VS^A + $dscp^C$ classifiers and dscp value matches.
75	VS^A + $Vlan^B$ + $dscp^C$	Profile with VS^A classifier. Profile with VS^A +IP or dot1d classifiers.
76	VS^A + $Vlan^B$ + $dscp^C$	Profile with VS^A + $Vlan^B$ + $dscp^C$ classifiers and dscp value matches.
77	VS^A + $Vlan^B$ + $dscp^C$	Profile with VS^A + $Vlan^B$ classifiers. Profile with VS^A + $Vlan^B$ + IP or dot1d classifiers.
78	VS^A + $Vlan^B$ + $dscp^C$	Profile with $Vlan^B$ +$dscp^C$ classifiers and $Vlan^B$ is a member of VS^A and dscp value matches.
79	VS^A + $Vlan^B$ + $dscp^C$	Profile with $Vlan^B$ classifier and $Vlan^B$ is a member of VS^A. Profile with $Vlan^B$ + IP or dot1d classifiers and $Vlan^B$ is a member of VS^A.
80	VS^A + $Vlan^B$ + $dscp^C$	Profile with IP or dot1d classifier.
81	VS^A + $Vlan^B$ + $dscp^C$	Profile with $dscp^C$ and dscp value matches.

Copyright© 2012-2013 Ciena® Corporation

Copyright© 2012-2013 Ciena® Corporation

39XX/51XX Service Delivery and Aggregation Switches

Configuration

Copyright© 2012-2013 Ciena® Corporation. All rights reserved.

Release 6.10
Publication: 009-3203-008
Document status: Standard
Revision F
Document release date: September 2013

CONTACT CIENA
For additional information, office locations, and phone numbers, please visit the Ciena web site at **www.ciena.com**

ciena