39XX/51XX Switches and Platforms

Advanced Ethernet Configuration

SAOS 6.21.2

What's inside...

New in this release
Advanced Ethernet configuration fundamentals
Access Control List
Hardware resource management
System timing configuration
VLAN management
MEF L2 VPNs
Quality of Service configuration
Multicast services configuration
TDM services over PSN configuration
Layer 2 Control Frame Tunneling configuration
Broadcast containment
Multi-chassis link aggregation
CFM-based link aggregation
Port state mirror group
Rapid Spanning Tree Protocol (RSTP) configuration
Multiple Spanning Tree Protocol (MSTP) configuration
Ethernet Ring Protection switching configuration
Virtual Link Loss Indication configuration
Service access control
Flow loop detection
Error codes

009-3458-040 - Standard Revision A
March 2022

Copyright© 2022 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. For the most up-to-date technical publications, visit www.ciena.com.

Copyright© 2022 Ciena® Corporation. All Rights Reserved

Use or disclosure of data contained in this document is subject to the Legal Notices and restrictions in this section and, unless governed by a valid license agreement signed between you and Ciena, the Licensing Agreement that follows.

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		*www.ciena.com/support/*
Sales and General Information	North America: 1-800-207-3714 International: +44 20 7012 5555	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	+91-22-42419600	E-mail: sales@ciena.com
In Latin America	011-5255-1719-0220 (Mexico City)	E-mail: sales@ciena.com
Training		E-mail: learning@ciena.com

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

READ THIS LICENSE AGREEMENT ("LICENSE") CAREFULLY BEFORE INSTALLING OR USING CIENA SOFTWARE OR DOCUMENTATION. THIS LICENSE IS AN AGREEMENT BETWEEN YOU AND CIENA COMMUNICATIONS, INC. (OR, AS APPLICABLE, SUCH OTHER CIENA CORPORATION AFFILIATE LICENSOR) ("CIENA") GOVERNING YOUR RIGHTS TO USE THE SOFTWARE. BY INSTALLING OR USING THE SOFTWARE, YOU ACKNOWLEDGE THAT YOU HAVE READ THIS LICENSE AND AGREE TO BE BOUND BY IT.

1. License Grant. Ciena may provide "Software" to you either (1) embedded within or running on a hardware product (together with Software, "Product") or (2) as a standalone application, and Software includes upgrades acquired by you from Ciena or a Ciena authorized reseller. The terms of this License apply to your use of the Software and associated documentation whether such has been provided by Ciena, an affiliate of Ciena, or by means of an authorized reseller or distributor. Subject to these terms, and payment of all applicable License fees including any usage-based fees, Ciena grants you, as end user, a non-exclusive, non-transferable, personal License to use the Software only in object code form, subject to any applicable authorized use, activation requirements, usage levels, scope of functionality and release level of the Software, as set forth in the applicable quote accepted by Buyer upon Buyer's issuance of an acceptable purchase order ("Order"), and in accordance with the detailed ordering information in the Ciena's generally available, applicable, Product documentation as of the date of such Order. Unless the context does not permit, Software also includes associated documentation. Where an Order is for a (a) perpetual license, you may use the Software and associated documentation for as long as you use the Product for internal business use, or a (b) subscription license, you may only use the Software and associated documentation during the subscription term. A subscription license includes Software upgrades and/or technical support Services during the subscription term (that are not included in a perpetual license), in accordance with the Order and as further described in the applicable Ciena's service description as of the date of the applicable Order. Prior to the expiration of each subscription term, Ciena will send you a quote for the annual renewal fee(s). To renew the subscription Software license(s) for additional subscription terms, you issue an Order in advance of the then-current expiration date of such subscription term.

2. Open Source and Third-Party Licenses. If any Software is subject to an open-source license that provides the end user with rights that are broader than this License, then such rights shall take precedence. Ciena warrants that using Software in accordance with its documentation will not subject you to any obligation to disclose, distribute or license your own software that interacts with Software.

3. Title. You are granted no title or ownership rights in or to the Software. Unless specifically authorized by Ciena in writing, you are not authorized to create any derivative works based upon the Software. Title to the Software, including any copies or derivative works based thereon, and to all copyrights, patents, trade secrets and other intellectual property rights in or to the Software, are and shall remain the property of Ciena and/or its licensors. Ciena's licensors are third party beneficiaries of this License. Ciena reserves to itself and its licensors all rights in the Software not expressly granted to you.

4. Confidentiality. The Software contains trade secrets of Ciena. Such trade secrets include, without limitation, the design, structure and logic of individual Software programs, their interactions with other portions of the Software, internal and external interfaces, and the programming techniques employed. The Software and related technical and commercial information, and other information received in connection with the purchase and use of the Software that a reasonable person would recognize as being confidential, are all confidential information of Ciena ("Confidential Information").

5. Obligations. You shall:

i) Hold the Software and Confidential Information in strict confidence for the benefit of Ciena using your best efforts to protect the Software and Confidential Information from unauthorized disclosure or use, and treat the Software and Confidential Information with the same degree of care as you do your own similar information, but no less than reasonable care;
ii) Keep a current record of the location of each copy of the Software you make;
iii) Use the Software only in accordance with the authorized usage level;
iv) Preserve intact any copyright, trademark, logo, legend or other notice of ownership on any original or copies of the Software, and affix to each copy of the Software you make, in the same form and location, a reproduction of the copyright notices, trademarks, and all other proprietary legends and/or logos appearing on the original copy of the Software delivered to you; and
v) Issue instructions to your authorized personnel to whom Software is disclosed, advising them of the confidential nature of the Software and provide them with a summary of the requirements of this License.

Copyright© 2022 Ciena® Corporation

6. Restrictions. You shall not:

i) Use the Software or Confidential Information a) for any purpose other than your own internal business purposes; and b) other than as expressly permitted by this License;
ii) Allow anyone other than your authorized personnel who need to use the Software in connection with your rights or obligations under this License to have access to the Software;
iii) Make any copies of the Software except such limited number of copies, in machine readable form only, as may be reasonably necessary for execution in accordance with the authorized usage level or for archival purposes only;
iv) Make any modifications, enhancements, adaptations, derivative works, or translations to or of the Software;
v) Reverse engineer, disassemble, reverse translate, decompile, or in any other manner decode the Software;
vi) Make full or partial copies of the associated documentation or other printed or machine-readable matter provided with the Software unless it was supplied by Ciena in a form intended for reproduction;
vii) Export or re-export the Software and/or the associated documentation from the country in which it was received from Ciena or its authorized reseller unless authorized by Ciena in writing; or
viii) Publish the results of any benchmark tests run on the Software.

7. Audit: Upon Ciena's reasonable request you shall permit Ciena to audit the use of the Software to ensure compliance with this License.

8. 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 Software and any accompanying technical data (collectively “Materials”) are commercial within the meaning of applicable Federal acquisition regulations. The Materials were developed fully at private expense. U.S. Government use of the Materials is restricted by this License, and all other U.S. Government use is prohibited. In accordance with FAR 12.212 and DFAR Supplement 227.7202, the Software is commercial computer software and the use of the Software is further restricted by this License.

9. Term of License. This License is effective until the applicable subscription term expires or the License is terminated. You may terminate this License by giving written notice to Ciena. This License will terminate immediately if (i) you breach any term or condition of this License or (ii) you become insolvent, cease to carry on business in the ordinary course, have a receiver appointed, enter into liquidation or bankruptcy, or any analogous process in your home country. Termination shall be without prejudice to any other rights or remedies Ciena may have. Upon any termination of this License, you shall destroy and erase all copies of the Software in your possession or control, and forward written certification to Ciena that all such copies of Software have been destroyed or erased. Your obligations to hold the Confidential Information in confidence, as provided in this License, shall survive the termination of this License.

10. Compliance with laws. You agree to comply with all laws related to your installation and use of the Software. Software is subject to U.S. export control laws and may be subject to export or import regulations in other countries. If Ciena authorizes you to import or export the Software in writing, you shall obtain all necessary licenses or permits and comply with all applicable laws.

11. Limitation of Liability. ANY LIABILITY OF CIENA SHALL BE LIMITED IN THE AGGREGATE TO THE AMOUNTS PAID BY YOU TO CIENA OR ITS AUTHORIZED RESELLER 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 LICENSOR OF CIENA. NEITHER CIENA NOR ANY OF ITS LICENSORS 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.

12. General. Ciena may assign this License to an affiliate or to a purchaser of the intellectual property rights in the Software. You shall not assign or transfer this License or any rights hereunder, and any attempt to do so will be void. This License shall be governed by the laws of the State of New York without regard to conflict of laws provisions. The U.N. Convention on Contracts for the International Sale of Goods shall not apply hereto. This License constitutes the complete and exclusive 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© 2022 Ciena® Corporation

Publication history

March 2022

Standard Revision A

First standard release of this document for SAOS 6.21.2.

Copyright© 2022 Ciena® Corporation

Copyright© 2022 Ciena® Corporation

Contents

About this document xxiii

New in this release 1-1

Revision A 1-1

Advanced Ethernet configuration fundamentals 2-1

Access control list 2-1
Hardware resources 2-2
System timing 2-2
Virtual Local Area Networks 2-2
MEF Layer 2 Virtual Private Networks 2-3
Layer 2 control frame tunneling 2-3
Quality of Service 2-6
Multicast services 2-6
TDM services over PSN configuration 2-7
Rapid Spanning Tree Protocol 2-8
Multiple Spanning Tree Protocol 2-9
Ethernet Ring Protection 2-10
- Logical rings 2-10
- Virtual rings 2-10
- Virtual sub-rings 2-11
- Ring protection process 2-11
- R-APS PDUs 2-13

Virtual Link Loss Indication 2-14
Service Access Control 2-14

Access Control List 3-1

Feature interaction 3-2
ACL resource usage 3-5
ACL device modes 3-5
ACL profile 3-6
ACL rules 3-7
Applying an ACL profile to an interface 3-10
ACL profile hierarchy 3-11
Feature limitations 3-11
Procedures 3-16

Copyright© 2022 Ciena® Corporation

List of procedures

3-1 Configuring global ACL modes 3-18
3-2 Creating an ACL profile 3-19
3-3 Adding or removing ACL rules to or from an ACL profile 3-20
3-4 Applying ACL profiles 3-23
3-5 Viewing ACL information 3-26
3-6 Clearing ACL rule hit statistics 3-30
3-7 Enabling and disabling an ACL profile 3-31
3-8 Enabling and disabling ACL 3-32

Hardware resource management 4-1

Virtual port multicast group resources 4-6

List of procedures

4-1 Configuring resources 4-8
4-2 Configuring hardware access control list resources 4-10
4-3 Freeing all hardware access control list resources 4-12
4-4 Configuring broadcast containment resources 4-13
4-5 Freeing all broadcast containment resources 4-15
4-6 Configuring CFM resources 4-16
4-7 Freeing all CFM resources 4-17
4-8 Configuring loss measurement resources 4-18
4-9 Freeing all loss measurement resources 4-20
4-10 Configuring traffic profiling resources 4-21
4-11 Freeing all traffic profiling resources 4-23
4-12 Configuring transport OAM resources 4-24
4-13 Freeing all transport OAM resources 4-26
4-14 Configuring virtual circuit statistics resources 4-27
4-15 Freeing all virtual circuit statistics resources 4-29
4-16 Configuring virtual circuit transform resources 4-30
4-17 Freeing all virtual circuit transform resources 4-32
4-18 Configuring virtual switch L2 enhanced transform resources 4-33
4-19 Freeing all virtual switch L2 enhanced transform resources 4-35
4-20 Displaying resource configuration information 4-36
4-21 Addressing classifier resource allocation too small for current configuration error 4-37
4-22 Displaying resource configuration in the configuration file 4-39
4-23 Displaying system resources 4-40

System timing configuration 5-1

Synchronous Ethernet 5-2
IEEE 1588 version 2 Precision Time Protocol 5-4
G.8275.1 profile support 5-4
External timing 5-10
BITS 5-10
GPS 5-11
Clock selection algorithm 5-11
Quality level value definition 5-13
Frequency, phase and time-of-day configuration rules 5-17

Copyright© 2022 Ciena® Corporation

Holdover interval 5-18
PTP clock type 5-18
Network configuration examples 5-19
Procedures 5-23

List of procedures

5-1 Enabling and disabling synchronization 5-25
5-2 Configuring synchronization 5-26
5-3 Configuring the global PTP telecom profile 5-29
5-4 Configuring the PTP timing global attributes 5-30
5-5 Configuring global attributes for PTP input timing 5-32
5-6 Configuring global attributes for PTP output timing 5-33
5-7 Configuring global attributes for GPS output timing 5-35
5-8 Configuring SyncE input references 5-36
5-9 Configuring BITS input references 5-38
5-10 Configuring PTP input references 5-41
5-11 Configuring GPS input references 5-44
5-12 Configuring SyncE output references 5-46
5-13 Configuring BITS output reference 5-47
5-14 Configuring PTP output timing references 5-49
5-15 Configuring GPS output references 5-51
5-16 Configuring protection-groups 5-52
5-17 Clearing timing statistics 5-60
5-18 Displaying information for synchronization 5-61
5-19 Displaying SyncE information 5-65
5-20 Displaying BITS information 5-68
5-21 Displaying PTP information 5-71
5-22 Displaying GPS information 5-78
5-23 Displaying TDM information 5-81
5-24 Displaying frequency information 5-83
5-25 Displaying phase information 5-85
5-26 Displaying time-of-day information 5-87
5-27 Displaying protection-group information 5-89
5-28 Sample configuration: system timing by means of SyncE and PTP 5-91
5-29 Sample configuration: system timing by means of PTP boundary clock 5-93

VLAN management 6-1

VLANs and traffic flow 6-1
Acceptable Frame Types 6-1
Port VLAN ID 6-2
VLAN ingress filter 6-2
Egress untagged VLAN 6-3
Behavior summary 6-3
VLAN/port configuration 6-4
VLAN translation 6-5
VLAN-based cross-connection 6-6
TPID configuration 6-7
Procedures 6-10

List of procedures

6-1 Changing the TPID stamp for a VLAN 6-11

Copyright© 2022 Ciena® Corporation

6-2 Configuring the boxwide inner TPID for a 39XX/51XX switch 6-12
6-3 Configuring a VLAN/port pair to ingress tagged traffic and egress tagged traffic 6-14
6-4 Configuring a VLAN/port pair to ingress untagged traffic and egress untagged traffic 6-16
6-5 Changing tag status 6-18
6-6 Configuring hybrid traffic 6-20
6-7 Emulating a tagged Ethernet port 6-22
6-8 Translating a single NNI VLAN 6-24
6-9 Translating a dual NNI VLAN 6-26
6-10 Translating a double-tagged VLAN 6-29
6-11 Creating a VLAN cross-connection 6-31
6-12 Deleting a VLAN cross-connection 6-32
6-13 Displaying VLAN cross-connections 6-33
6-14 Cross connection deployment example 6-34

MEF L2 VPNs 7-1

Overview 7-1
Q-in-Q encapsulation 7-2
EPL and EVPL Provider Bridge configuration 7-4
Provider VLAN 7-5
Reserved VLAN 7-6
Virtual switch 7-7
Virtual circuits 7-8
Frame flooding behavior 7-9
Q-in-Q Ethertype 7-9
Virtual switch emulation of a VLAN on a UNI port 7-10
Translate-tag zero and double tagged UNI frames 7-13
UNI member statistics 7-13
Traffic profiles 7-14
Statistics and statistics resources 7-16
Feature interaction 7-18
Caveats 7-18
Use case 7-19
EVPL CoS 7-20
Virtual switch CoS policies 7-20
Virtual switch member CoS policy override 7-20
EVPL Bundling 7-21
CVID range configuration 7-22
Virtual switch SVLAN transforms 7-25
Virtual switch L2 transforms 7-27
Ingress push; egress pop (i-push,e-pop) 7-28
Ingress push; egress pop and stamp (i-push,e-pop:stamp) 7-28
Ingress stamp; egress match, pop, and stamp (i-stamp:push,e-match-pop:stamp) 7-29
Untagged broadcast, unknown unicast and multicast traffic enhancements 7-33
Virtual circuit transforms 7-34
Configuration model 7-36
VC transform creation 7-38
Applying a VC transform to a VS 7-40

Copyright© 2022 Ciena® Corporation

Applying a VC transform to a VS cross-connect 7-40
CFM on a VC transform service configuration 7-40
Resource manager 7-41
VC transform feature interactions 7-41
VC transform considerations 7-44
Private forwarding groups 7-45
L2-based PFG Forwarding policy sets 7-47
Port-based PFG egress profile 7-49
Port membership 7-52
Feature interaction 7-53
Upgrading and downgrading a device 7-54
External Network-to-Network Interface Hairpin 7-54
Sub-port interfaces 7-55
MEF services 7-58
E-Line 7-58
E-LAN 7-59
E-Tree 7-60
E-Access 7-60
MPLS encapsulation of MEF8 TDM PWs 7-61

List of procedures

7-1 Creating an EPL provider bridge 7-62
7-2 Creating an EVPL provider bridge 7-64
7-3 Configuring a virtual switch to emulate a VLAN on a UNI port 7-66
7-4 Configuring the fixed encapsulation priority value 7-69
7-5 Configuring UNI member statistics 7-71
7-6 Clearing UNI member statistics 7-75
7-7 Displaying UNI member information 7-76
7-8 Determining egress resources to collect Tx statistics 7-78
7-9 Setting the CoS policy when adding VS members 7-79
7-10 Handling ingress untagged frames 7-81
7-11 Setting the virtual switch L2 transform action on a port 7-83
7-12 Creating an i-push, e-pop: stamp configuration 7-84
7-13 Creating an i-stamp:push,e-match-pop:stamp configuration 7-88
7-14 Configuring flood traffic from a NNI port toward a UNI port 7-91
7-15 Configuring a virtual circuit transform 7-94
7-16 Configuring VC transform support on the 3928 and 5160 switches 7-95
7-17 Configuring a virtual circuit transform service 7-97
7-18 Configuring a remote management interface on a VC VLAN on a VC transform 7-100
7-19 Configuring a virtual circuit transform to a virtual circuit 7-103
7-20 Removing a virtual circuit transform from a virtual circuit 7-104
7-21 Setting the virtual circuit transform to default values 7-105
7-22 Enabling or disabling a virtual circuit transform 7-106
7-23 Deleting a virtual circuit transform 7-107
7-24 Displaying virtual circuit transform 7-108
7-25 Displaying system-reserved port information for the 5160 switch 7-109
7-26 Preventing NNI ingress frames without P-Tag from being switched 7-111
7-27 Configuring L2 PFGs 7-113
7-28 Configuring port-based PFGs 7-117

Copyright© 2022 Ciena® Corporation

7-29 Configuring egress profile with traffic types 7-121
7-30 Disabling the PFG feature 7-124
7-31 Displaying the configuration of EVPL VS members 7-125
7-32 Displaying virtual switches 7-126
7-33 Displaying PFG information 7-128
7-34 Configuring EPL and EVPL for E-Access service types 7-130
7-35 Configuring CFM for E-Access service types 7-134
7-36 Configuring HIM for E-Access service types 7-137
7-37 Configuring L2 Control Frame Tunneling for E-Access service types 7-139
7-38 Configuring RFC 2544 Benchmarking for E-Access 7-140
7-39 Configuring ENNI hairpin switching using sub-ports 7-145
7-40 Displaying statistics for sub-port interfaces 7-148
7-41 Clearing statistics for sub-port interfaces 7-150
7-42 Configuring a virtual switch for remote management and customer data traffic transiting a device 7-151
7-43 Creating an i-push, e-pop Q-in-Q VS configuration 7-156
7-44 Configuring MPLS encapsulation of MEF8 TDM PWs 7-158
7-45 Configuring the CVID range as a single EVPL 7-160
7-46 Setting the VLAN range for a port 7-162
7-47 Unsetting the VLAN range for a port 7-163
7-48 Displaying the VLAN range for a device 7-164
7-49 Configuring a Q-in-Q virtual switch with SVLAN transform 7-165
7-50 Displaying information related to Q-in-Q virtual switch with SVLAN transform 7-168

Quality of Service configuration 8-1

Class of Services (CoS) policies and mapping 8-1
RCoS and RColor behavior for QinQ virtual switch 8-4
Ports with virtual switch members 8-5
Remarking behavior for VLAN switching 8-6
RCoS to FCoS mappings for Ethernet networks 8-9
Traffic profiling 8-10
Global profile mode 8-10
Global traffic profiles 8-11
Determining classification attributes 8-13
Determining metering attributes 8-17
Traffic profiling with VLAN range support 8-20
Traffic profiling with hierarchical ingress metering 8-21
Advanced classification mode 8-22
392X Traffic profile scale 8-23
Standard Traffic Profile (Port Mode) 8-24
Resource manager 8-27
MEF 10.3 ingress metering 8-27
Envelope 8-30
Bandwidth profile parameters 8-30
Classifiers and actions 8-33
Use case 1 8-35
Use case 2 8-37
Use case 3 8-39
Use case 4 8-40

Copyright© 2022 Ciena® Corporation

Use case 5 8-42
Envelope statistics 8-44
Interaction with standard and global traffic profiles 8-45
Limitations 8-46
Congestion management 8-48
Creating and modifying sWRED profiles on the 3903, 3903x, 3904, 3905, 3924, 3926, 3928, 3942, 5142, and 5160 8-48
Egress queuing and scheduling 8-50
Non-hierarchical egress shaping 8-50
Hierarchical egress queuing 8-54
Local CPU-generated control, management and OAM traffic QoS treatment 8-61
DSCP and PCP assignment per protocol 8-62
Determining RCoS from DSCP or PCP 8-64
RCoS to egress queue mapping 8-65
RCoS to FCoS mapping at egress port 8-66
Remote management interface CoS behavior 8-66
Transit node behavior for DSCP and PCP 8-67
Request or reply type frames DSCP and PCP behavior 8-68
RCoS and egress queue assignment for untagged traffic 8-72
Configuration example 8-72
Limitations 8-72
Dynamic microflow detection and greedy flow monitoring 8-73
Basic working 8-76
General feature interactions 8-81
Limitations 8-82
Recommended ports 8-82
Configuration 8-82
Devices supported 8-86
Procedures 8-87

List of procedures

8-1 Configuring Class of Services (CoS) policies on a port 8-88
8-2 Displaying the CoS policy and mapping on a port 8-90
8-3 Creating a resolved CoS map 8-91
8-4 Modifying a resolved CoS map 8-93
8-5 Setting the resolved CoS map for a port 8-94
8-6 Deleting a custom RCoS map 8-95
8-7 Displaying resolved CoS maps 8-96
8-8 Creating a frame CoS map 8-98
8-9 Modifying a frame CoS map 8-100
8-10 Deleting a frame CoS map 8-101
8-11 Enabling and disabling the remarking of CPU generated frames 8-102
8-12 Setting the frame CoS map for a port 8-103
8-13 Displaying frame CoS maps 8-104
8-14 Configuring ingress RCoS to egress queue mapping 8-105
8-15 Displaying an RCoS map 8-107
8-16 Applying RCoS policies and mapping in a VLAN 8-108
8-17 Applying RCoS policies and mapping in a virtual switch 8-112
8-18 Configuring an L2-only type RCoS to FCoS map 8-115
8-19 Configuring a global type RCoS to FCoS map 8-116

Copyright© 2022 Ciena® Corporation

8-20 Displaying the RCoS to FCoS map 8-117
8-21 Enabling and disabling standard traffic profiling 8-118
8-22 Enabling and disabling standard traffic extended resources 8-119
8-23 Configuring an envelope 8-120
8-24 Configuring a bandwidth profile flow in an envelope 8-121
8-25 Clearing statistics for envelope flows 8-123
8-26 Deleting an envelope 8-124
8-27 Enabling an envelope flow 8-125
8-28 Disabling an envelope flow 8-126
8-29 Removing an envelope flow 8-127
8-30 Modifying bandwidth flow attributes 8-128
8-31 Displaying envelopes 8-129
8-32 Setting default values for envelope and flow attributes 8-132
8-33 Configuration example for MEF 10.3 ingress metering 8-133
8-34 Configuration example of a c-pcp type envelope 8-136
8-35 Configuration example of a vs-only type envelope 8-138
8-36 Enabling and disabling global profiling 8-140
8-37 Enabling and disabling global parent profiling 8-141
8-38 Setting the traffic profiling provisioning mode, profile mode, and profile type 8-142
8-39 Displaying traffic profiling information in port profile mode 8-144
8-40 Setting traffic profiling port attributes for standard profiles 8-147
8-41 Configuring a traffic profiling standard profile 8-149
8-42 Configuring a traffic profiling global traffic profile 8-155
8-43 Configuring traffic profiling global profile parent 8-161
8-44 Configuring per-port standard traffic profiling 8-163
8-45 Configuring per-port and per-VLAN standard traffic profiling 8-166
8-46 Configuring hierarchical VLAN standard traffic profiles 8-169
8-47 Configuring hierarchical standard port traffic profiles 8-175
8-48 Configuring VS classification for standard traffic profiles 8-177
8-49 Setting the VLAN range for standard traffic profiles 8-179
8-50 Configuring standard traffic profiling for EVPLs with the vlan-range-type as group 8-182
8-51 Configuring a global profile with the VLAN range option 8-184
8-52 Configuring VS classification for standard HIM traffic profiles 8-187
8-53 Displaying standard traffic profiles 8-189
8-54 Clearing statistics for all standard profiles 8-191
8-55 Clearing statistics for all global profiles 8-192
8-56 Clearing statistics for all global parent profiles 8-193
8-57 Displaying statistics for all standard profiles 8-194
8-58 Displaying statistics for all global profiles 8-195
8-59 Displaying statistics for global parent profiles 8-196
8-60 Displaying attributes for all global profiles 8-197
8-61 Displaying attributes for global parent profiles 8-199
8-62 Displaying throughput statistics for a global profile 8-200
8-63 Displaying throughput statistics for a global parent profile 8-201
8-64 Displaying throughput statistics for a standard traffic profile 8-202
8-65 Setting the per-port hierarchical traffic-profiling mode 8-203
8-66 Setting the child mode 8-204
8-67 Creating an sRED profile 8-205

Copyright© 2022 Ciena® Corporation

8-68 Creating an sWRED profile 8-206
8-69 Modifying an sWRED profile 8-207
8-70 Creating an sWRED profile 8-208
8-71 Modifying an sWRED profile 8-209
8-72 Displaying custom congestion avoidance profiles 8-210
8-73 Updating the congestion avoidance profile for an egress port queue 8-212
8-74 Clearing the congestion avoidance profile to the default for an egress port queue 8-213
8-75 Displaying the congestion avoidance profile for an egress port queue 8-214
8-76 Deleting a custom congestion avoidance profile 8-215
8-77 Renaming custom congestion avoidance profile 8-216
8-78 Changing the algorithm of an egress port scheduler for non-hierarchical egress queuing 8-217
8-79 Changing the weight of the scheduler for a queue 8-218
8-80 Displaying queue weight and scheduler algorithms for non-hierarchical egress queuing 8-219
8-81 Configuring egress port and queue shaping for non-hierarchical egress queuing 8-220
8-82 Displaying egress port queue configuration for non-hierarchical egress queuing 8-223
8-83 Displaying egress port queue statistics for non-hierarchical egress queuing 8-225
8-84 Configuring hierarchical egress port shaping 8-227
8-85 Configuring the default queue group for hierarchical egress queuing 8-228
8-86 Configuring the queue group profile for hierarchical egress queuing 8-230
8-87 Allocating the number of queue groups for each physical port on the 3924, 3926 and 3928 8-232
8-88 Deallocating the number of queue groups for each physical port on the 3924, 3926 and 3928 8-234
8-89 Displaying the number of queue groups for each physical port on the 3924, 3926 and 3928 8-235
8-90 Configuring the dynamic queue group profile for hierarchical egress queuing 8-236
8-91 Configuring the custom queue group for hierarchical egress queuing 8-238
8-92 Adding a port to a custom queue group 8-239
8-93 Associating a custom queue group with a virtual circuit for hierarchical egress queuing 8-241
8-94 Associating a custom queue group with a VLAN for hierarchical egress queuing 8-245
8-95 Configuring the egress queue depth for hierarchical egress queuing 8-249
8-96 Configuring the port scheduler for hierarchical egress 8-250
8-97 Displaying hierarchical egress queuing information 8-251
8-98 Configuring frame bandwidth calculation 8-256
8-99 Sample configuration for local CPU-generated control, management and OAM traffic QoS treatment 8-258
8-100 Configuring a device for dynamic microflow detection and greedy flow monitoring 8-260
8-101 Minimally configuring microflow policing on port 1 8-263
8-102 Working examples of dynamic microflow detection and greedy flow monitoring 8-264

Copyright© 2022 Ciena® Corporation

Multicast services configuration 9-1

IGMP snooping 9-3
Enhanced features 9-4
Multicast operations 9-5
Multicast forwarding domains 9-6
Multicast interface 9-6
Multicast traffic filters 9-6
Multicast servers and routers 9-7
Server topology 9-8
IGMP query engine 9-9
Router IP address range 9-9
Channel stream 9-10
Rapid recovery mode 9-11
Layer 2 layer 3 mismatch behavior 9-11
IGMP groups 9-12
Statistics 9-12
IGMPv3 9-14
IGMPv3 messages 9-14
IGMP snooping 9-15
Membership reporting 9-15
Interface state 9-16
No host suppression 9-18
Queries 9-18
Services supported 9-21
Limitations 9-21
IGMP interworking 9-22
Interface state change 9-23
Port state change 9-23
RSTP and MSTP events 9-24
G.8032 9-24
LACP 9-28
MC-LAG 9-28
MPLS 9-30
Multicast-services attributes 9-31

List of procedures
9-1 Configuring a VLAN as a multicast L2 forwarding domain 9-34
9-2 Configuring a multicast service on a virtual switch 9-35
9-3 Configuring channel streams 9-36
9-4 Configuring IGMP forking with VLAN translation 9-38
9-5 Configuring a multicast router topology 9-40
9-6 Configuring a multicast server topology 9-44
9-7 Configuring multicast servers with redundant routers 9-47
9-8 Configuring redundant query engines 9-51
9-9 Clearing multicast service statistics 9-55
9-10 Displaying multicast services information 9-56
9-11 MPLS-based multicast configuration example 9-58

TDM services over PSN configuration 10-1

Copyright© 2022 Ciena® Corporation

Ports 10-2
TDM SFP ports 10-3
TDM module ports 10-6
Combo TDM OC48 FRU module ports 10-6
Port alarms 10-7
Signal degrade and signal fail 10-7
Protection groups 10-8
Encapsulation 10-9
Transport 10-10

E1 cable pin configuration 10-10
E1 channel number mapping 10-11
T1 SFP pin mapping 10-13
TDM FRU pin mapping 10-14
DS1 channel numbers 10-14
Linear APS 1+1 10-17
APS 1+1 on the Combo TDM OC48 FRU module 10-18
APS 1+1 feature restrictions 10-20
Linear APS modes of operation 10-21
Linear APS switchover criteria 10-23
Linear APS switch commands 10-24
APS status 10-26
APS protocol signaling and K bytes 10-26
Linear APS 1+1 caveats and limitations 10-30
Framing 10-31
Channel Associated Signaling (CAS) for nxDSO service 10-31
nxDS0 CAS-transport configuration notes 10-32
Feature limitations 10-32
Ethernet feature restrictions 10-33
Section trace 10-33
Pseudo-Random Bit Sequence 10-33
Double tags 10-34
MPLS encapsulation 10-34
Timing 10-35
Packet delay variation 10-35
Transceiver-related limitations 10-36
Procedures 10-39

List of procedures

10-1 Configuring the FRU type 10-43
10-2 Configuring TDM services 10-45
10-3 Creating TDM SFP ports 10-48
10-4 Deleting TDM SFP ports 10-54
10-5 Loading a new firmware load on an SFP 10-57
10-6 Creating and enabling a port to be used by a TSOP transceiver 10-61
10-7 Changing the section trace 10-63
10-8 Configuring a protection group 10-64
10-9 Configuring a line-code on a T1 module 10-65
10-10 Configuring a line-code on an E1 module 10-66
10-11 Configuring framing on a T1 module 10-67

Copyright© 2022 Ciena® Corporation

10-12 Configuring framing on a E1 module 10-68
10-13 Configuring framing on a Combo TDM OC48 FRU module 10-69
10-14 Configuring SSM 10-71
10-15 Configuring loopcode transmission on T1 module ports 10-73
10-16 Configuring C-TPID on the 3926 T1/E1 port 10-75
10-17 Configuring C-TPID on the TDM SFP 10-76
10-18 Configuring the cable length 10-77
10-19 Configuring the line build out cable length 10-78
10-20 Configuring TDM SFP PRBS insertion and monitoring 10-79
10-21 Recovering the clock 10-85
10-22 Configuring TDM global synchronization 10-91
10-23 Configuring TDM synchronization frequency protection groups 10-92
10-24 Clearing TDM port statistics 10-95
10-25 Configuring TDM profiles 10-96
10-26 Configuring attachment circuits 10-100
10-27 Configuring signal degrade 10-104
10-28 Configuring signal fail 10-105
10-29 Configuring OC3 port loopbacks 10-106
10-30 Configuring OC12 port loopbacks 10-109
10-31 Configuring OC48 port loopbacks 10-112
10-32 Configuring DS3 port loopbacks 10-114
10-33 Configuring E1 port loopbacks 10-115
10-34 Configuring T1 port loopbacks 10-116
10-35 Configuring TDM virtual circuits 10-118
10-36 Configuring MPLS transport 10-136
10-37 Configuring Ethernet virtual circuits 10-137
10-38 Configuring virtual switch cross-connections 10-138
10-39 Configuring jitter buffer monitoring on a T1 module 10-139
10-40 Displaying information about the FRU module in 3926 10-142
10-41 Displaying FPGA information about FRU modules 10-143
10-42 Displaying TDM port information 10-145
10-43 Displaying attachment circuit information 10-152
10-44 Displaying TDM port loopbacks 10-153
10-45 Displaying virtual circuit information 10-156
10-46 Displaying virtual switch information 10-158
10-47 Displaying virtual switch cross-connection information 10-159
10-48 Configuring SAToP services over 802.1Q Metro Ethernet on a TDM SFP or FRU port 10-160
10-49 Configuring CESoP services over 802.1Q Metro Ethernet on a TDM SFP or FRU port 10-162
10-50 Configuring SAToP services over Q-in-Q Metro Ethernet on a a TDM SFP or FRU port 10-164
10-51 Configuring SAToP pseudowire over an MPLS network 10-167
10-52 Configuring 1+1 linear APS groups 10-170
10-53 Adding TDM ports to 1+1 linear APS groups 10-172
10-54 Enabling 1+1 linear APS groups 10-175
10-55 Switching over 1+1 linear APS groups 10-178
10-56 Showing and clearing 1+1 linear APS group statistics 10-180

Layer 2 Control Frame Tunneling configuration 11-1

Copyright© 2022 Ciena® Corporation

L2 control frame tunneling profile 11-9
Tunnel method 11-9
Transparent mode 11-11
L2PT mode 11-12
Tagged L2CF handling 11-13
Configuration 11-15
Profile processing 11-15
L2-CFT port protocol conflicts 11-16

List of procedures

11-1 Adding and removing untagged L2 control frame classification 11-19
11-2 Configuring control-frame tunneling profiles 11-20
11-3 Enabling and disabling L2 control frame tunneling 11-22
11-4 Enabling and disabling tagged Cisco PVST L2PT forwarding 11-23
11-5 Configuring the destination MAC address for L2PT tunnel method 11-24
11-6 Adding and removing control protocols 11-25
11-7 Configuring the disposition of control protocols 11-27
11-8 Setting L2 control frame tunneling attributes 11-28
11-9 Displaying enabled L2 control frame tunneling instances 11-29
11-10 Displaying tagged Cisco PVST L2PT configured values 11-30
11-11 Displaying L2-CFT resources 11-31
11-12 Displaying the L2 control frame classification for a port 11-33
11-13 Displaying L2 control frame tunneling configuration in the configuration file 11-35
11-14 Configuring L2 control frame tunneling 11-37
11-15 Configuring L2 control frame tunneling mode for MEF CE 2.0 compliance 11-44

Broadcast containment 12-1

Feature interactions 12-3
Scale management 12-4
Non-resource mode 12-4
Resource mode 12-4
Limitations 12-6

List of procedures

12-1 Creating a broadcast containment filter 12-7
12-2 Setting broadcast containment filter attributes 12-9
12-3 Setting broadcast containment configuration 12-11
12-4 Enabling and disabling broadcast containment 12-12
12-5 Deleting a broadcast containment filter 12-13
12-6 Adding ports to a broadcast containment filter 12-14
12-7 Removing a port from a broadcast containment filter 12-15
12-8 Displaying broadcast containment filters 12-16
12-9 Displaying a broadcast containment filter to which per port per VS membership is attached 12-22
12-10 Displaying broadcast containment configuration in the configuration file 12-23
12-11 Displaying all broadcast containment resources 12-24
12-12 Configuring broadcast containment filters on a port and classifier combination of port and virtual switch 12-25

Copyright© 2022 Ciena® Corporation

12-13 Clearing broadcast containment statistics 12-27

Multi-chassis link aggregation 13-1

Hashing 13-2
Monitoring inter-chassis link connections 13-2
Detecting failures 13-3
MC-LAG interoperability between SAOS 6.x and 8.x 13-3
Procedures 13-3
13-1 Configuring a multi-chassis link aggregation group 13-4
13-2 Adding a port to a multi-chassis link aggregation group 13-6
13-3 Adding a protection port to a multi-chassis link aggregation group 13-7
13-4 Forcing a switchover on a redundancy group node 13-8
13-5 Configuring an inter-chassis link 13-9
13-6 Displaying an inter-chassis link 13-10
13-7 Displaying a summary of MC-LAGs 13-11

CFM-based link aggregation 14-1

14-1 Configuring CFM service fault-based LAG switchover 14-4

Port state mirror group 15-1

Port mirroring 15-1
Mirror frame encapsulation 15-2
Egress port mirroring behavior 15-2
Port state mirroring 15-3
Uni-directional port state mirroring 15-4
Bi-directional port state mirroring 15-5
Limitations 15-6
Ingress mirroring and maximum frame size 15-7
Egress mirroring and maximum frame size 15-7

List of procedures
15-1 Configuring port mirroring 15-8
15-2 Configuring encapsulated port mirroring 15-9
15-3 Configuring port state mirroring groups 15-11

Rapid Spanning Tree Protocol (RSTP) configuration 16-1

Broadcast storms 16-4
STP/RSTP goals 16-5
Rapid Spanning Tree Protocol implementation 16-5
Bridge Protocol Data Units 16-5
Electing the root 16-7
Links and path costs 16-8
Port states and roles 16-9
Protocol timers 16-11
Converging the network with RSTP 16-12
Topology changes 16-12
Spanning tree domains 16-15
Loop detection and prevention 16-16
RSTP standards compliance 16-17

Copyright© 2022 Ciena® Corporation

List of procedures

16-1 Changing RSTP standards compliance 16-18
16-2 CLI configuration example 16-19

Multiple Spanning Tree Protocol (MSTP) configuration 17-1

Overview 17-1
MST regions and spanning trees 17-2
Bridges 17-2
BPDUs 17-2
Port roles 17-5
Boundary port 17-6
Network convergence 17-7
Backward compatibility 17-7
Configuration 17-9
Configuring the CIST root bridge 17-10

List of procedures

17-1 Configuring MSTP CIST attributes 17-12
17-2 Configuring CIST ports 17-13
17-3 Configuring the MST Identifier 17-15
17-4 Changing mode 17-17
17-5 Configuring MSTIs 17-19
17-6 Changing the MSTI bridge priority 17-23
17-7 Deleting an MSTI 17-25
17-8 Configuring MSTI ports 17-26
17-9 Configuring bridge global settings 17-28
17-10 Configuring ports 17-30
17-11 Enabling and disabling MSTP on specified ports 17-32
17-12 Enabling MSTP and saving configuration 17-33
17-13 Disabling and enabling MSTP on a bridge 17-34
17-14 Disabling and enabling MSTP on an MSTI 17-35
17-15 Displaying MSTP information 17-36
17-16 Clearing MSTP statistics 17-39

Ethernet Ring Protection switching configuration 18-1

MC-LAG interworking with G.8032 18-2
Procedures 18-2

List of procedures

18-1 Verifying that ring protection is enabled 18-4
18-2 Disabling ring protection 18-5
18-3 Configuring logical rings 18-6
18-4 Configuring virtual rings 18-13
18-5 Verifying ERP switching operation 18-21
18-6 Displaying information for logical rings 18-22
18-7 Displaying virtual ring information 18-24
18-8 Displaying ring protection 18-27
18-9 Displaying ERP switching in the configuration file 18-28
18-10 Checking configuration for a ring not forwarding traffic 18-29
18-11 Configuration example 18-30

Copyright© 2022 Ciena® Corporation

Virtual Link Loss Indication configuration 19-1

VLLI signaling 19-1
VLLI architecture 19-3
Y-cable protection use case 19-5
Y-cable protection failure and reversion scenarios 19-6
Procedures 19-7

List of procedures
19-1 Configuring VLLI to shut down a port when CCM loss is detected - virtual switch 19-8
19-2 Configuring VLLI to shut down a port when CCM loss is detected - VLAN 19-13
19-3 Configuring VLLI between a UNI and an ENNI 19-16
19-4 Configuring inverse VLLI with one 39XX/51XX switch and two PEs 19-20
19-5 Configuring inverse VLLI with two 39XX/51XX switches and two PEs 19-23
19-6 Configuring VLLI for MAC flush on an intersecting node 19-26
19-7 Configuring MAC flush from the MPLS network ingress node to all other nodes 19-28
19-8 Y-cable client protection using inverse VLLI between nodes B and A 19-30
19-9 MPLS active/active protection state management using VLLI in a Y-cable topology 19-33
19-10 MPLS active/standby protection state management using VLLI in a Y-cable topology 19-42
19-11 Configuring VLLI to block a pseudowire 19-49
19-12 Generating VLLI Fault and Recovery events based on pseudowires 19-51

Service access control 20-1

List of procedures
20-1 Configuring SAC with MAC aging 20-3
20-2 Configuring static MACs and SAC on the same device 20-6
20-3 Configuring the device to forward or drop unknown frames 20-10

Flow loop detection 21-1

Loop detection 21-3
Loop notification 21-4
Loop prevention 21-5

List of procedures
21-1 Enabling and disabling flow loop detection 21-6
21-2 Enabling MAC learning at a virtual switch or VLAN 21-7
21-3 Configuring flow loop detection 21-8
21-4 Resetting flow loop detection to default values 21-10
21-5 Clearing loop detection statistics on the specified port(s) 21-11
21-6 Displaying flow learning information 21-12
21-7 Displaying flow loop detection information 21-13

Error codes 22-1

Traffic profiling error codes 22-1

Copyright© 2022 Ciena® Corporation

About this document

This document describes how to configure advanced Ethernet features on 39XX/51XX Switches and Platforms. 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.

To configure Advanced Ethernet features, you need to install the Advanced Ethernet license key. License keys can be purchased by contacting Ciena Customer Support.

> ***Note 1:*** System timing commands are disabled if the AE and AOAM license key is not installed.
>
> ***Note 2:*** This system software cannot be installed on any other Service Delivery Switches, Service Concentration Switches or Service Aggregation Switches.

This document 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, these text conventions are used:

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

Copyright© 2022 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. This 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 can be entered as server 10.10.11.100 or server www.ciena.com. With priority <NUMBER: 1-7> the text in <> 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© 2022 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.

Documents in the 39XX/51XX documentation suite

For descriptions of documents in the 39XX/51XX Switches and Platforms documentation suite, see *39XX/51XX Switches and Platforms Product Fundamentals*.

Copyright© 2022 Ciena® Corporation

Copyright© 2022 Ciena® Corporation

New in this release

Revision A

This section provides a summary of the changes made in *39XX/51XX Switches and Platforms Advanced Ethernet Configuration* for SAOS 6.21.2.

In this chapter "Multi-chassis link aggregation", we made updates relating to the feature to add more than one primary or protection ports as MC-LAG members.

Changes to:

- "Configuring a multi-chassis link aggregation group" on page 13-4.
- "Adding a port to a multi-chassis link aggregation group" on page 13-6.
- "Adding a protection port to a multi-chassis link aggregation group" on page 13-7.
- In the chapter "Virtual Link Loss Indication configuration", we added a new feature to extend and monitor support to pseudowires in an N:1 and N:M relationship, during a fault or recovery scenario in a VLLI configuration "Generating VLLI Fault and Recovery events based on pseudowires".
- In the chapter Chapter 5, "System timing configuration", the system timing support is updated to include the 3924 platform.

Copyright© 2022 Ciena® Corporation

Copyright© 2022 Ciena® Corporation

Advanced Ethernet configuration fundamentals

This section provides an overview of the components and protocols required to configure Advanced Ethernet:

- “Access control list”
- “Hardware resources”
- “System timing”
- “Virtual Local Area Networks”
- “MEF Layer 2 Virtual Private Networks”
- “Layer 2 control frame tunneling”
- “Quality of Service”
- “Multicast services”
- “TDM services over PSN configuration”
- “Rapid Spanning Tree Protocol”
- “Multiple Spanning Tree Protocol”
- “Ethernet Ring Protection”
- “Virtual Link Loss Indication”
- “Service Access Control”

Access control list

A hardware Access Control List (ACL) is an ordered set of rules that is used to filter ingress/egress traffic on a device. This set of rules is applied directly to an interface and contains a default rule, or action, that is applicable if none of the other rules applied to the interface are matched. All non-default rules contain a set of layer 2, layer 3 and/or layer 4 packet filter terms, as well as a target action. The action of a particular rule is executed if all filter terms in the rule have been matched. The supported packet actions either “allow” or “deny” the packet from being switched. All packets which ingress an interface are processed by the ACL associated with the interface when the ACL feature is

Copyright© 2022 Ciena® Corporation

enabled. Incoming packets get tested against a particular hierarchy of rules based on the interface (such as port, VLAN, virtual switch) the ingress frame resolves to.

Hardware resources

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

Customized resource management is supported for these features:

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

System timing

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

System timing comprises:

- Synchronous Ethernet (SyncE)
- external interfaces, which comprise:
  - Building Integrated Timing Supply (BITS)
  - Global Positioning System (GPS)
- Time Division Multiplexing (TDM) line timing
- Pseudowire emulation (PWE)

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 to limit the distribution of unicast, multicast, and broadcast traffic. For example, flooded traffic originating from a particular VLAN can be limited to only other ports

Copyright© 2022 Ciena® Corporation

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 in 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 VPNs" on page 7-1.

MEF Layer 2 Virtual Private Networks

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

- MEF 6.2 services: Associating two more UNIs in a service provider Carrier Ethernet Network
  - E-Line: point-to-point service
  - E-LAN: multipoint-to-multipoint service
  - E-Tree: rooted multipoint service
- MEF 51 OVC services: in an operator Carrier Ethernet Network
  - E-Access: point-to-point and multipoint service with UNIs and ENNIs
  - E-Transit: point-to-point and multipoint service with ENNIs

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.

Layer 2 control frame tunneling

Layer 2 control frame tunneling controls the processing of Layer 2 control frames (L2CF) associated with various control protocols. This control applies to L2CFs ingressing on UNI interfaces. When an untagged L2CF is received by the switching chip, it is either sent to the CPU to be handled by a protocol manager process, or it is discarded. L2CFT provides a way to modify this default behavior so that the L2CFs for certain protocols can be transformed and forwarded as data frames, optionally replacing the standard protocol MAC-DA with a configurable Layer 2 protocol tunneling (L2PT) MAC DA. L2CFT also handles the transformation of L2PT tunneled frames back into its standard protocol form.

Copyright© 2022 Ciena® Corporation

When a tagged L2CF is received by the switch, it is forwarded as data frames. L2CFT tagged handling provides a mechanism to override this default behavior. It operates in service unaware mode.

Service unaware tagged L2CFT handling applies tagged L2CFT knobs on all tagged frames ingressing on a port. This functionality is supported on UNI interfaces. Depending on the state of processing, L2 control frames can be in three forms, as described in this table.

Table 2-1
Layer 2 control frames

L2 control frames	Description
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 on 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. Layer 2 control frames may be received on the subscriber facing port in this form. Additionally, this form may occur when an untagged control frame is received from a subscriber facing interface, encapsulated as a data frame, and then forwarded as a tagged frame from a network facing interface for transport through the provider network. Transparent and tagged frames are the same. Tagged L2CFT has the capability to handle transparently tagged control frames on UNI frames.

Copyright© 2022 Ciena® Corporation

Table 2-1
Layer 2 control frames

L2 control frames	Description
L2 Protocol Tunneling (L2PT)	Standard MAC DA has been transformed to the L2PT special MAC DA. The frame's original Mac DA is replaced with a configurable L2PT MAC address. The L2PT DA MAC can be configured with a valid Multicast MAC address. The default system L2PT MAC is the Generic Bridging PDU Tunneling (GBPT) MAC of 01:00:0C:CD:CD:CD:D0. It is not recommended to use L2PT MAC belonging to these MAC address blocks: • L2 Slow Protocol Block: 01:80:C2:XX:XX:XX • ISO 9542 Address: 09:00:2B:00:00:04/05 • IEEE802.5 Block: 03:00:00:XX:XX:XX • IPv4 Multicast Block: 0.1:00:5E:XX:XX:XX • IPv6 Multicast Block: 33:33:XX:XX:XX:XX

Copyright© 2022 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. This table lists mechanisms for managing bandwidth.

Table 2-2
Mechanisms for managing bandwidth

Mechanism	Description
Class of Services (CoS) policies and mapping	CoS policies and mapping comprise: • Resolved CoS: applies a Resolved CoS Policy and Resolved CoS Map on ingress. Optionally, remarks the frame's Layer 2 priority and color based on the mapping. Assigns traffic to egress CoS queues based on 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. • 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

Copyright© 2022 Ciena® Corporation

implementation was designed for networks where multicast services are critical, such as networks delivering IP broadcast video. Devices employ enhanced IGMP snooping and various filters to limit multicast streams and assure their timely delivery.

These network elements support 1023 multicast groups:

- 3903
- 3903x
- 3904
- 3905
- 3906
- 3924
- 3926
- 3928
- 3942
- 5142
- 5160

TDM services over PSN configuration

Time Division Multiplexing (TDM) services allow the network provider to create TDM connections through a Packet Switched Network (PSN).TDM services are available on these switches:

- 3904
- 3905
- 3924
- 3926
- 3928
- 5142
- 5160

Every port on these switches can support TDM SFP. TDM SFP ports on these switches provide TDM services.

TDM services are:

- Structure-Agnostic TDM over Packet (SAToP)
- Circuit Emulation Services over Packet Switched Networks (CESoP)

Copyright© 2022 Ciena® Corporation

SAToP TDM pseudowire takes the raw TDM bit-stream from T1 or E1, disregards any TDM structure imposed on it, prepends the pseudowire header and dispatches the resulting packet stream over the Ethernet or MPLS PSN. At the remote end, the pseudowire header information is used to identify the egress attachment circuit. The TDM payload from the pseudowire is played out over the egress attachment circuit using the clock information as defined when the pseudowire is configured. The PWE hardware supports up to 16 pseudowire services for SAToP whether the configuration is ANSI or ETSI.

Note: SAToP TDM services are only supported on PWE module ports.

With CESoP, structured TDM signals are encapsulated in a pseudowire and sent across the PSN. CESoP services are supported on TDM SFP ports and PWE module ports. Advantages of CESoP are:

- PSN bandwidth savings
- DS0 level grooming
- better resiliency to packet loss in the PSN
- more flexible distribution

Rapid Spanning Tree Protocol

Spanning Tree Protocol (STP) is a Layer 2 link management protocol defined in IEEE 802.1D. STP provides a loop-free topology in a bridged network and delivers efficient reconfiguration of the loop-free topology in the event that a link fails. Spanning Tree allows devices in an Ethernet network to dynamically create loop-free paths by configuring ports that create loops into a standby (blocking) state.

ATTENTION
Rapid Spanning Tree Protocol is the default and preferred xSFP protocol.

While link redundancy is desirable for increased fault tolerance, a network topology that has loops (more than one active path to reach end stations) can result in inaccurate forwarding, inaccurate learning, and proliferation of broadcast messages.

Rapid Spanning Tree Protocol (RSTP), formerly defined in IEEE 802.1w and now incorporated in IEEE 802.1D, was developed to achieve faster convergence times. In most circumstances, RSTP can converge the network in less than two seconds. This is critical for networks that carry delay-sensitive traffic such as voice and video. RSTP achieves faster convergence times by selecting not only primary paths to the root bridge, but also backup paths. RSTP is backwards compatible with equipment using the IEEE 802.1D spanning tree algorithm standard.

Copyright© 2022 Ciena® Corporation

A primary goal in a Ciena Access Network is to provide failure protection in the 50 ms time frame. This time requirement makes Ciena Carrier Grade Ethernet solutions competitive with SONET failover times as well as meeting the time requirements necessary to protect voice circuits from dropping calls. RSTP reconvergence (change of topology) times can often be below 50 ms.

- Fault notification - Fault notification is provided by the MEP that detected a connectivity fault for one of these reasons:
  - expected CCMs were not received
  - unexpected or invalid CCM were received
  - the CCM carried a notification of the failure of its associated MEP
- Fault recovery - Fault notifications help network operators correct configuration errors or replace failed components.

Multiple Spanning Tree Protocol

Multiple Spanning Tree Protocol (MSTP) provides:

- Scalability on large networks. Both STP and RSTP 802.1W allow the total BPDU traversing of maximum 24 bridges from the root bridge. RSTP 802.1D/D4 (2003) allows the BPDU traversing of maximum 40 bridges. In an MST region, MSTP limits BPDU traversing to a maximum of 30 bridges, with no limit on the number of regions. In the case of a ring topology, the number of bridges a BPDU can traverse is limited to the maximum number of bridges configured on the ring, taking the worst case of link failure into consideration. In a meshed topology, the maximum number of bridges that can be included in the topology is much more than the limit.
- Limited Topology changes. MSTP limits the topology change events in the region as much as possible.
- Redundant links and load balancing. MSTP supports the assignment of a number of VLANs to a single spanning tree instance, and multiple instances are supported in each spanning tree region. These multiple Spanning Tree instances enable L2 networks to have redundant physical links because VLANs can be assigned to different virtual bridges (spanning tree instances); this also enables load balancing.

ATTENTION

Rapid Spanning Tree Protocol is the default and preferred xSTP protocol. If MSTP is not required, but configured, change the xSTP mode to RSTP, save the configuration and reboot the device.

ATTENTION

If MSTP is the desired xSTP protocol, already configured and going to be used in conjunction with G.8032 Ring Protection, the same VLAN cannot exist in both the G.8032 and MSTP realms.

Copyright© 2022 Ciena® Corporation

Ethernet Ring Protection

Ethernet Ring Protection (ERP) switching, as defined by the G.8032 standard, protects the connectivity between network nodes using the Ring Automatic Protection Switching (R-APS) control protocol. Each node is connected to two adjacent nodes by means of a physical port or link aggregation port forming a closed loop.

With ERP switching, you can configure a set of interconnection nodes to operate as if connected in a ring or multiple rings as shown in this figure.

Figure 2-1
Interconnection nodes configured as a ring

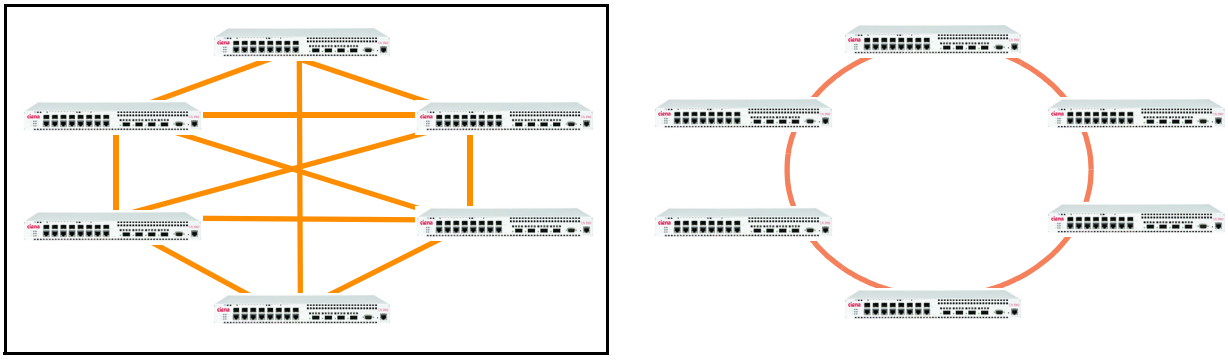

> **ATTENTION**
> If MSTP is already configured on the device, confirm that MSTP is the required xSTP protocol. If MSTP is not required, change the configured xSFP mode to RSTP, save the configuration and reboot the device.

> **ATTENTION**
> The same VLAN cannot exist in both the G.8032 and MSTP realm.

Logical rings

The logical ring defines the connections between the nodes. Connections are identified by setting a logical west port and a logical east port. The ring configuration for each device represents a segment of the ring.

Virtual rings

Multiple virtual rings can be configured against a single logical ring to carry different types of traffic that can be segregated by assigning one or more VLANs to a virtual ring. Each virtual ring is managed by its own R-APS channel which exchanges messages over a dedicated VLAN. These messages include communication of virtual ring topology changes, operator command notification, such as force switching, and general virtual ring status.

Copyright© 2022 Ciena® Corporation

Note: The data channel for transporting client service traffic is defined by the configuration of enterprise VLANs or E-LINE, E-TREE, and E-LAN.

Virtual sub-rings

Virtual rings support a multi-ring or ladder network with the configuration of a virtual sub-ring connected to a major virtual ring. The sub-ring has its own R-APS channel and contains its messages to the sub-ring network by terminating messages at the interconnection node.

A virtual sub-ring defined with east or west port termination is considered to be an interconnection node to another major virtual ring or virtual sub-ring. That is, the defined termination port is managed by the virtual sub-ring while ring ports on the interconnection link are managed by the connected major virtual ring or virtual sub-ring.

This figure shows interconnection nodes configured in a major virtual ring and virtual sub-ring.

Figure 2-2
Interconnection nodes configured in a major virtual ring and virtual sub-ring

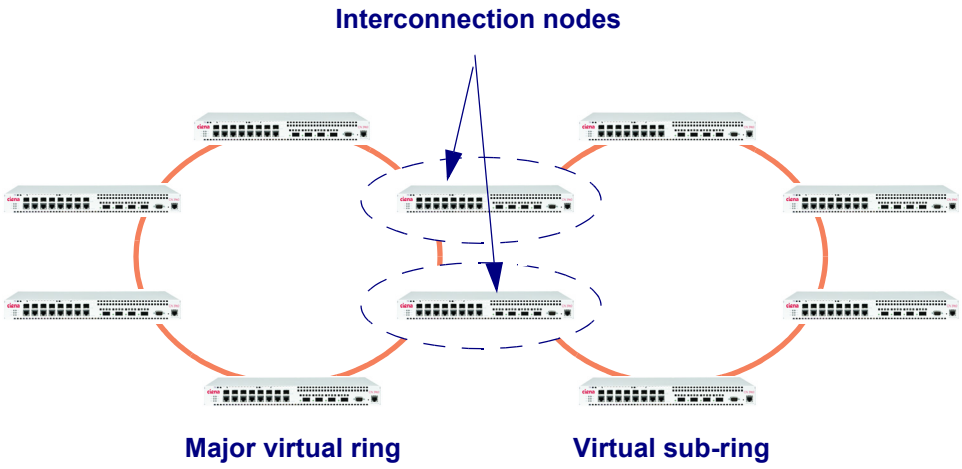

Ring protection process

Each virtual ring must have one port designated as a channel block called the Ring Protection Link (RPL). When a fault occurs, the ports adjacent to the failure become blocked, and transmit a R-APS failure (signal failure) message to adjacent and downstream nodes. Downstream nodes that receive an R-APS failure message flush their forwarding database. If the downstream node is the RPL owner, the RPL channel block is also removed, allowing traffic to flow over the RPL port.

Copyright© 2022 Ciena® Corporation

On recovery of the failure, the blocked adjacent ports send R-APS recovery (node recovery) messages, and the guard timer is started to prevent reception of R-APS messages during transmission. Once the original channel block port receives the R-APS recovery messages, it waits a period of time defined by the wait-to-restore timer. After the wait-to-restore timer expires, the channel block is restored and sends a R-APS Okay (node recovered remove block) message. As the nodes receive the Okay message, the forwarding database is flushed, and the ports adjacent to the failure point remove their blocks restoring the ring to its original configuration.

> ***Note:*** R-APS sends all new messages at three millisecond intervals three times in rapid succession and then sends all subsequent messages at five second intervals.

In the case of the virtual sub-ring, the behavior is similar with a few differences:

- R-APS messages received at the sub-ring termination nodes terminate at the port.
- R-APS messages generated at the CPU are sent to the sub-ring's terminating port on terminated nodes and to both ports on regular nodes.
- R-APS messages received at non-terminating sub-ring ports are forwarded to the other ring port even if the port is in blocked state.
- R-APS messages received at each sub-ring port are delivered to the CPU.

Copyright© 2022 Ciena® Corporation

R-APS PDUs

R-APS messages are in the form of a R-APS Protocol Data Unit (PDU) containing the information shown in this table.

Table 2-3
R-APS PDU information

R-APS messages	Meaning
MEL	MEG level of the R-APS PDU
Version	Y.1731 Version field value of 0x01.
OpCode	Y.1731 Op
Flags	Y.1731 Flag field value of 0x00. This field is ignored for received messages.
TVL Offset	Y.1731 TLV Offset field value of 0x20 (32 decimal)
R-APS Specific Information	Contains these fields: • Request/State: Contains a value to represent the request or state, including Forced Switch, Event, Signal Failure, Manual Switch, or No Request. • Sub-code: Value is always transmitted as “0000.” If the Request/State field is Event, the sub-code value represents Flush Request. Otherwise, the sub-code is ignored for received messages. • Status Field: Value to show whether the RPL is blocked, whether the forwarding database is flushed, and reference to the blocked port (ring 0 or 1). • Nod ID: Chassis MAC address identifier unique to the node.
Optional TLV	Y.1731 Optional TLV. If left unspecified, it is replaced with the End TLV.
End TLV	Y.1731 End TLV field value of 0x00.

Copyright© 2022 Ciena® Corporation

The R-APS PDU format is shown in this table.

Table 2-4
R-APS PDU format

1								2								3								4							
8	7	6	5	4	3	3	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1	8	7	6	5	4	3	2	1
MEL			Version					OpCode								Flags								TLV Offset							
R-APS Specific Information																															
Optional TLV																								End TLV							

Virtual Link Loss Indication

Virtual Link Loss Indication (VLLI) allows network elements at the UNI-N of an EVC to signal to the UNI-C of the customer premise equipment (CPE) if:

- The remote UNI attachment interface has failed or recovered.
- The EVC connection has failed or recovered and impacting end-to-end client traffic.

VLLI applies to EVCs with a point-to-point topology.

Service Access Control

Service Access Control allows for static MAC addresses. It also controls the number of dynamically learned MAC addresses on a per-port basis.

By configuring static MAC entries, security is provided to allow only those frames to be forwarded that have a source MAC address that is recognized by the device.

Once the maximum number of dynamically learned MAC addresses is reached, additional MAC addresses are not learned and frames from subsequent devices can be configured to either be forwarded or dropped. The learn limit does not apply to static MAC addresses.

Copyright© 2022 Ciena® Corporation

Access Control List

A hardware Access Control List (ACL) is an ordered set of rules that is used to filter ingress/egress traffic on a device. This set of rules is applied directly to an interface and contains a default rule, or action, that is applicable if none of the other rules applied to the interface are matched. All non-default rules contain a set of layer 2, layer 3 and/or layer 4 packet filter terms, as well as a target action. The action of a particular rule is only executed if all filter terms in the rule have been matched. The supported packet actions either “allow” or “deny” the packet from being switched. All packets which ingress an interface are processed by the ACL associated with the interface when the ACL feature is enabled. Incoming packets get tested against a particular hierarchy of rules based on the interface (such as port, VLAN, virtual switch) the ingress frame resolves to.

ACLs enhance network security by allowing traffic from trusted sources, and protecting networks or hosts from unwanted or malicious traffic. ACLs can be used to track unauthorized access attempts, and prevent denial of service attacks and control traffic flow.

Note: Hardware ACL requires the Advanced Ethernet license.

Copyright© 2022 Ciena® Corporation

Feature interaction

Some SAOS features have special interactions with ACLs. Generally, any feature which is configured to drop packets always takes precedence over ACL allow rules. Similarly, ACL deny rules take precedence over forwarding behaviors of other features. This table outlines some, but not all, of these features.

Table 3-1
Feature interaction

Feature interaction	Behavior
Benchmark/RFC 2544	ACL rules can be assigned to the benchmark port under test; they are not applied to traffic locally generated or reflected by Benchmark.
Broadcast Containment	Frames discarded by broadcast containment filters are always discarded regardless of any configured ACL “allow” rules. Frames dropped by ACL “deny” rules still consume broadcast containment meter credits when broadcast containment has been configured with resource mode on. ACL statistics increment even if the frame was dropped by broadcast containment, and vice versa.
IP-ACL	ACL and IP-ACL can be configured simultaneously; ACL rules configured by means of the hardware are invoked before any rules configured by the IP-ACL feature. An exception to the above applied with ACL rules is applied to the local interface. In this scenario, the IP-ACL rules are executed before ACL rules.
L2 Control Frame Tunneling	A frame forward using L2 control frame tunneling is discarded if an ACL “deny” rule exists and qualifies on the same frame. Frames dropped by an L2 control frame tunneling “discard” disposition are always discarded regardless of any configured “allow” ACL rules. ACL statistics increment even if the frame was dropped by the L2 frame tunneling “discard” disposition.
Loop Prevention Protocols	A port that discards traffic due to a Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP) or G.8032 blocked state always discards frames regardless of any configured “allow” ACL rules. ***Note:*** All packets which ingress the device and match an ACL rule increment ACL statistics regardless of the action of other SAOS features.

Copyright© 2022 Ciena® Corporation

Table 3-1
Feature interaction

Feature interaction	Behavior
Port Mirroring	An ingress packet continues to be mirrored regardless of any configured ACL deny rule. An egress packet is not mirrored if the packet matches an ACL deny rule.
Private Forwarding Group	Traffic that is being discarded due to a configured Private Forwarding Group (PFG) policy are not forwarded even if an ACL “allow” rule exists. ACL statistics increment even if the frame was dropped by PFG forwarding rules.
Traffic Profiling	Frames discarded by traffic profiling are always discarded regardless of any configured ACL “allow” rules. Frames which are dropped by ACL “deny” rules still consume traffic profiling meter credits. ACL statistics always increment even if the frame was dropped by traffic profiling and vice-versa.

Copyright© 2022 Ciena® Corporation

Table 3-1
Feature interaction

Feature interaction	Behavior
TWAMP	Hardware ACL rules take precedence over any TWAMP access deny configurations.
VC/sub-port statistics	VC/sub-port statistics increment even if the traffic flow matches an ACL deny rule resulting in a dropped frame.
CFM/Y.1731	ACL can be configured to match and drop ingress CFM/Y.1731 frames on all non-accelerated services and sessions, that is. delay measurements, loopback, linktrace, synthetic frame loss and frame loss. ACL “deny” rules may match and increment statistics for accelerated CFM service Continuity Check Messages; these frames are still processed by the hardware. The state of the CFM service is not affected. Loopback messages received on an accelerated CFM service may match and increment ACL “deny” rules statistics; Loopback Reply messages are still sent by the hardware. Linktrace messages received on an accelerated CFM service are dropped if they match an ACL “deny” rule.’ The ACL feature can be configured to prevent processing of Delay Measurement, Synthetic Frame Loss and Frame Loss Messages for accelerated sessions on the 3903, 3904, 3905, 3906, 3924, 3926, 3928, 3942, and 5142 devices. On the 5160 device, these frames for accelerated sessions are processed before ACL rules. Any configured ACL rules are not hit. When the ACL feature is provisioned to drop a particular traffic flow, configured accelerated Loss Measurement sessions may still increment hardware counters even though some or all of the target data frames were not forwarded by the switch. In this case, the frame loss statistics reported by the target session are misleading.

In addition, these SAOS features have similar frame filtering capabilities as the ACL. These features may apply a drop action on specific traffic flows prior to an ACL rule enforcement. Since statistics for the ACL feature operate independently, the same traffic feature increments applicable ACL filter statistics. The current list of these features is

- VLAN translation (to discard frames)
- Service Access Control

ACL processing occurs for every packet that ingresses the switch. This results in ACL statistics being incremented regardless of the forwarding or drop decisions made by other SAOS features. For packet filtering, an ACL rule is

Copyright© 2022 Ciena® Corporation

only matched if the ingress packet matches the packet filtering criteria defined by the rule. Since ACL rules are associated with interfaces, a particular rule is only matched if the ingress packet is switched on the resolved interface the rule is associated with. The interface an ACL filter becomes associated with is equivalent to the resolved interface (port, vlan, virtual switch, IP interfaces) after the port PVID, L2 lookup, VLAN translation and other switching domain decisions occur.

ACL comprises these interactions:

- A drop of discard action in any one SAOS feature results in a dropped frame regardless of a forwarding action configured in the ACL feature, and vice versa.
- A drop or discard action in the ACL feature does not prevent statistics increments in the ACL feature for the same traffic flow.
- A drop or discard action in the ACL feature does not prevent meter consumption in a separate SAOS feature from the same traffic flow.

ACL resource usage

Using ACLs requires allocating hardware resources which are not allocated by default. The maximum number of ACL rules that can be applied in hardware is limited by the hardware resource availability in the resource manager. Provided that the maximum number of hardware resources is configured for the hardware ACL feature, then this table shows the maximum number of hardware rules that can be configured on each device.

Table 3-2
ACL hardware allocation

SAOS device	Maximum number of ACL rules
3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928	768
3942	1664
5142 and 5160	3584

ACL device modes

ACL has two device modes:

- L2/L3 combo mode
- L3 only mode

The global ACL device mode can be changed without a device reboot; no ACL profiles can exist on the device before making this change. The default device mode is L2/L3 combo mode. It is only possible to switch between modes if no profiles have been configured.

Copyright© 2022 Ciena® Corporation

This table describes the ACL device modes.

Table 3-3
ACL device modes

ACL device mode	Description
L2/L3 combo mode	ACL rules are created with all available L2, L3 and L4 packet classification parameters with the exception of Source and Destination IPv6 addresses. This mode provides greater flexibility because of the different packet classification parameters. This mode is the default ACL device mode.
L3 only mode	ACL rules can be created with only the available L3 and L4 packet classification parameters. Accepts IPv4 and IPv6 addresses filter terms. Does not accept L2 filter terms. ***Note:*** L3 only mode blocks all layer 2 protocols, including ARP, when they end with a deny any statement.

ACL profile

An ACL profile is a named collection of ACL rules that can be applied to one or more L2 or L3 interfaces. L2 logical interfaces, with respect to ACLs, are ports (physical and aggregate), VLANs, and virtual switches. L3 logical interfaces are IP interfaces, local interfaces and remote interfaces.

An ACL becomes active or operationally-enabled when applied to one or more interfaces. A profile instance can be defined as a combination of the ACL profile and a single interface that profile is attached to. A particular ACL profile can have many profile instances configured on the switch. As such, each rule in the profile can be associated with many interfaces. The separation of ACL creation from the association of applicable interfaces means that rules can be defined and modified before becoming operationally enabled.

To create an ACL profile, two mandatory parameters are required:

- Rule Name
- Default Filter Action (allow or deny)

Copyright© 2022 Ciena® Corporation

The default filter action represents a default rule that is applicable to that particular profile. That is, if no other rules in the profile are matched, then the action of this default rule is executed. This default rule has the lowest priority/precedence in the profile. For each ACL profile created, there is always a minimum of one ACL rule that exists; this is represented by the default filter action.

Applying an ACL profile with no rules and a default filter action of deny to a logical interface denies all traffic on that interface. It is the user's responsibility to properly configure the rules in the ACL profile before applying it to any interface.

It is recommended to always set the default filter action to deny, while specifying allow rules for all types of traffic that are expected to ingress the logical interface. This allows users to take full advance of the ACL and enhances the security/robustness of the switch and the network.

Note: All newly created ACL profiles are enabled by default.

ACL rules

Each ACL rule comprises an "ANDed" set of filter terms or packet classification parameters, and an associated filter action. One or more ACL rules make up a profile. The maximum number of rules in a profile list is 256 (including that profile's default rule). In the ACL profile, a frame can only match one ACL rule. The frame takes on the action of the matched ACL rule with the highest precedence.

The Maximum number of supported ACL profiles/rules per device table shows the maximum allowable number of ACL profiles and the maximum number of ACL rules that can be created for each device. The ACL rule numbers depicted in this table illustrate the maximum number of CLI configured rules.

Table 3-4
Maximum number of supported ACL profiles/rules per device

Platform	Maximum supported ACL profiles	Maximum number of ACL rules
3903, 3904, 3905, 3906, 3924, 3926, 3928	512	1024
3942	1024	2048

Adding rules to an ACL

Copyright© 2022 Ciena® Corporation

An ACL rule must specify a number of CLI attributes. These attributes must be entered in this order:

- Rule Identifier (string)
- Rule Precedence value which determines the rule order in the ACL
- Filter action (allow or deny) to perform when a packet matches the filter terms
- Set of L2/L3/L4 filter terms or wildcard filter term (match all traffic)

This table lists mandatory ACL rule items.

Table 3-5
Mandatory ACL rule items

Mandatory item	Description
Rule Identifier (string)	A unique rule name specified by the user when the rule is created. Is used to modify or remove rules in a profile.
Rule Precedence	A unique numeric value 0-255 that must be specified during the creation. ***Note:*** The precedence value 255 is reserved for a profile's default rule. Indicates the priority or ordering of the rule in the ACL profile, 0 is the highest priority and 255 is the lowest priority. This attribute can be modified after the rule has been created.
Filter Action	An action that is to be executed if the rule filter terms are matched. Valid actions include "allow" and "deny".

Copyright© 2022 Ciena® Corporation

Mandatory item	Description
Filter Terms	A list of packet matching criteria. These L2 filter terms can be specified when adding an ACL rule to a profile: • Source MAC address with mask • Destination MAC address with mask • Base Ethertype • PCP with mask • DEI bit • VLAN ID with mask These L3/L4 filter terms can be specified when adding an ACL rule to a profile: • Source IPv4/v6 address with CIDR mask • Destination IPv4/v6 address with CIDR mask • IP protocol • IP fragment • DSCP with mask • L4 source port (single/range) • L4 destination port (single/range) • L4 application port (single/range) • L4 application (Telnet, SSH, TWAMP, HTTP, etc.) • TCP flags (SYNC, ACK, RST, URG, PSH, FIN) ***Note 1:*** Only the L4 destination port or L4 application filter term can be specified in a rule. ***Note 2:*** If the IP protocol is specified, it must be set to TCP or UDP when the L4 source or L4 destination port filter terms are used. ***Note 3:*** To classify on all traffic, the wildcard filter term “any” can be specified to match all traffic. ***Note 4:*** The wildcard filter term cannot be specified in any other L2 or L3 filter terms. If the “any” filter term is not used, at least one L2 or L3 filter term must be specified. Configuring ACL deny rules with generic filter terms may result in the discarding of key protocol frames such as ARP, CFM, LACP, LLDP, and G.8032.

Statistics are available on a per-rule basis as well as a per-rule-per-interface basis. These statistics are enabled by default.

Copyright© 2022 Ciena® Corporation

Applying an ACL profile to an interface

After the ACL profile is created and the rules defined, it next has to be applied to a logical interface. When an ACL is applied to at least one logical interface, it is considered active or enabled. An ACL that is not applied to any logical interfaces is considered not active. This table describes the logical interfaces that an ACL profile can be applied to.

Table 3-6
Logical interfaces

Logical interface	Description
Port	ACL rules applied to a port are applicable to all traffic that ingresses that physical or aggregated port. Any L2, L3 and/or L4 filter terms can be used on rules applied to this interface type.
VLAN	ACL rules applied to a VLAN are applicable to all traffic that ingresses all members of that VLAN and is equivalent to applying the same ACL to each port member of the VLAN, with the exception of statistics. ACL rules that are applied to a VLAN increment for ingress traffic on all members of that VLAN. Any L2, L3 and/or L4 filter terms can be used on rules applied to this interface type. New port members added to a VLAN are automatically subject to the ACL rules applied to the VLAN. ***Note:*** The VID filter term that is part of an ACL rule refers to the “off-the-wire” VLAN ID in the packet. The VLAN that an ACL profile applies to is the switching domain. Typically, the incoming VID off the wire is the same as the switching domain VLAN. In some situations, the VLAN translation may be used and the incoming VID may be different from the VLAN the frame is switched on.
Virtual Switch	ACL rules applied to a virtual switch are applicable to all traffic that ingresses virtual switch UNI members. Any L2, L3 and/or L4 filter terms can be used on rules applied to t his interface type. ***Note:*** ACL rules can be applied to a particular Ethernet Virtual Private LAN (EVPL) virtual switch member, such as port 1 and VID 100, by specifying the VID filter term to be 100 and applying the ACL to port 1.

Copyright© 2022 Ciena® Corporation

Table 3-6
Logical interfaces

Logical interface	Description
L3 interface	L3 interfaces include the IP interfaces, the local and remote interface. ACL rules apply to all terminating frames which ingress an IP interface, the remote interface or the local interface. ***Note:*** L3 interfaces control traffic reaching the CPU. Any L2, L3 and/or L4 filter terms can be used on rules applied to this interface type, with the exception of the local interface which only accepts L3 and L4 filter terms.

ACL profile hierarchy

A frame looks to match rules in a PORT ACL profile before moving onto other ACL profiles, such as VLAN, VS and L3 interfaces. The ACL profile hierarchy is as follows:

- Port ACL profile is the first set of ACL rules that a packet matches on. This profile has the highest priority.
- VLAN/VS ACL profiles exist at the same level. Note ACL restrictions prevent users from configuring ACL profiles on the active VLAN of VS. This means that a frame never matches both a VLAN and a VS rule.
- L3 interface ACL profile is the last set of ACL rules that a packet matches on. This profile has the lowest priority.

These conditions outline the ACL rule matching behavior:

- A non-default rule match results in the frame not being processed by any other ACL rules configured on the device.
- A default “deny” rule hit automatically results in a dropped packet. This means that the frame is not processed by other matching ACL rules on the device.
- A default allow rule hit causes the frame to be forwarded onto the next ACL profile in the hierarchy (VLAN, VS, L3 interface) for processing.

Feature limitations

This table describes the ACL hardware feature limitations.

Copyright© 2022 Ciena® Corporation

Table 3-7
Feature limitations

Limitation	Description
ACL rules on the local interface	ACL rules support a subset of the available L3/L4 filter terms: • Source IPv4/IPv6 address • Destination IPv4/IPv6 address • IP protocol • L4 source/destination port (single values only) • L4 application and TCP flags Any rules configured by means of IP-ACL that impacts traffic ingressing the local interface are hit before any rules in an ACL profile that is applied to the local interface. Traffic may be dropped by IP-ACL and not further processed by the ACL process.
Changing attributes of an active ACL rule	Changing any attribute of an ACL rule that is part of an active ACL profile, meaning one that is already applied to a logical interface, may cause the ACL rule to become ineffective for a small number of transiting frames.
Fragmented IP traffic and L4 filter terms	If a fragmented packet passes through an ACL with rules containing any L4 filter terms (L4 source/destination port and/or TCP flags), those rules only apply to the initial fragment or head of the packet. They do not apply to the subsequent fragments because the initial fragment is the only fragment containing L4 information.
Broadcast/multicast traffic on the ACL profiles associated with the remote or IP interfaces	ACL rules that are applied to an IP interface or the remote only apply to unicast traffic. Broadcast/multicast traffic do not match rules on an ACL profile applied on an IP or the remote interface. To filter on specific multicast or broadcast traffic, then those ACL rules need to be applied on a different logical interface that is associated with the L3 interface (VLAN/port) depending on the services configured. This restriction does not apply to ACL rules on the local interface.

Copyright© 2022 Ciena® Corporation

Limitation	Description
L4 source and destination port ranges	Filtering on L4 source and destination port ranges is only supported in L2/L3 combo mode. Single L4 source and/or destination port classification is allowed in both L2/L3 combo mode, as well as L3 only mode. A maximum of eight unique L4 source/destination port ranges can be specified for all ACL rules configured on the switch. Both L4 source port range and L4 destination port range filter terms cannot be specified with the same rule. There is no restriction for the single value L4 source or destination port filter terms.
IPv6 addresses	ACL rules with IPv6 address filter terms are supported in L3 device mode only.
Maximum number of ACL rules per active profile	ACL supports up to 255 user configurable rules per ACL profile, excluding the default rule. There is no guarantee that this number can be supported on all active ACL profiles configured on a device. The number of active ACL rules a user can configure on the device largely depends on the hardware resource availability on the device.
Frames dropped by ACL deny rules	Unicast frames normally learned by the L2 MAC table are always learned regardless of any configured ACL deny rules. This may result in undesirable behavior when discarded ACL traffic consumes addresses in the MAC table which are counted against the maximum number of addresses that the device can learn.

Copyright© 2022 Ciena® Corporation

Limitation	Description
MPLS and Tripled Tagged Frames	ACL L3/L4 filtering is limited on these encapsulation types: • MPLS — For terminating MPLS frames on a label edge router (LER), the L2 filter terms vtag-stack, Src and Dst MAC addresses are supported. All other filter terms are not supported. The “vtag-stack” filter term always matches on the outer VLAN in the outer L2 header. The Source and Destination MAC filter terms match on MAC addresses in the inner L2 header. For transiting frames being switched by an LSR, or L2 switch, ACL rules only qualify on the L2 transport header portion of MPLS frames. All remaining filter terms are not officially supported for MPLS terminating frames. The L2 filter terms are only matched as specified by ACL rules. • Tripled Tagged — SAOS 6.x devices do not detect triple-tagged frames. This type of frame can be switched by a device; only the first two tags are recognized. These frames are unable to match any L3/L4 fields that are present after the third VLAN tag.

Copyright© 2022 Ciena® Corporation

Limitation	Description
ACL filter terms for port types	When ACL rules are defined on Q-in-Q VS member ports or Q-in-Q VSes: • Ingress frames which have two or more VLAN tags have these filter term restrictions: — Base Ether Type qualifies on the TPID field of the second customer VLAN tag. — ACL rules do not match on the encapsulated L3/L4 data. • Ingress frames which have an outer TPID value that is not 0x8x8100 have these filter term restrictions: — Base Ether Type qualifies on the TPID field of the first invalid customer VLAN tag. — ACL rules do not match on the encapsulated L3/L4 data. When ACL rules are defined on non-Q-in-Q VS member ports or non-Q-in-Q VSes: • Ingress frames which have three or more VLAN tags have these filter term restrictions: — Base Ether Type qualifies on the TPID field of the third VLAN tag. — ACL rules do not match on the encapsulated L3/L4 data. • Ingress frames which have two VLAN tags where the inter tag TPID does not equal 0x8100 have these filter term restrictions: — Base Ether Type qualifies on the TPID field of the second invalid VLAN tag. — ACL rules do not match on the encapsulated L3/L4 data. • Ingress frames which have an outer VLAN tag TPID value that is not 0x8100, 0x9100, or 0x88a8 have these filter term restrictions: — Base EtherType qualifies on the TPID field of the first invalid VLAN tag. — ACL rules do not match on the encapsulated L3/L4 data.

Copyright© 2022 Ciena® Corporation

Limitation	Description
IP protocol filtering on IPv6 traffic with extension headers	If an ACL rule is configured with an IP protocol filter term and IPv6 traffic is expected to hit this rule, then the IP protocol filter term matches a different field in the packet depending on the presence/absence of IPv6 extension headers as outlined: • The IPv6 packet does not contain any extension headers. The IP protocol filter term matches on the “next header” value in the IPv6 header. • The IPv6 packet contains one or more IPv6 extension headers. The IP protocol filter term matches on the “next header” value in the first extension header after the IPv6 header.
Applying ACL profiles to VLANs	An ACL profile can only be applied to VLANs which are not used as the active VLAN of a virtual switch. Specifically, an ACL profile cannot be applied to a virtual circuit VLAN or a reserved VLAN.
Applying ACL profiles to VS member IP interfaces	ACL profiles cannot be directly applied to a VS member IP interface (EVC Ping IP interface). For traffic destined to a VS member IP interface, apply the ACL profile on the virtual switch that the VS member IP interface is a member of. To further classify the traffic destined to the EVC Ping IP interface, the VID filter term can be used and set to the CVID of the EVC ping traffic for each rule in the ACL profile.
Frames dropped by frame integrity and ACLs on L3 interfaces	Due to a hardware restriction, frames that are dropped by the Frame Integrity feature (address invalid frames, TCP malformed frames or ICMP oversized frames) do not increment ACL rule statistics for ACL rules that are applied on an IP interface or the Remote Interface. ***Note:*** This restriction does not apply to the local interface since the Frame Integrity feature does not operate on that interface. ACL rules on ports, VLANs and virtual switches still increment statistics for frames that are dropped by the Frame Integrity feature.

Procedures

This figure shows the ordered flow of ACL configuration procedures.

Copyright© 2022 Ciena® Corporation

Figure 3-1
ACL configuration

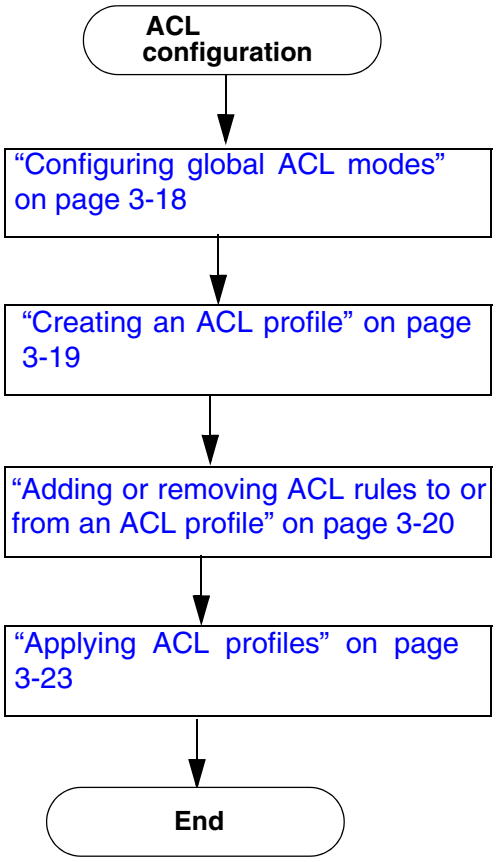

This section provides these procedures for ACL:

- "Configuring global ACL modes" on page 3-18
- "Creating an ACL profile" on page 3-19
- "Adding or removing ACL rules to or from an ACL profile" on page 3-20
- "Applying ACL profiles" on page 3-23
- "Viewing ACL information" on page 3-26
- "Clearing ACL rule hit statistics" on page 3-30
- "Enabling and disabling an ACL profile" on page 3-31
- "Enabling and disabling ACL" on page 3-32

Copyright© 2022 Ciena® Corporation

Procedure 3-1
Configuring global ACL modes

You can configure

- L2/L3 combo mode (default)
- L3 only mode
- ACL mode

  Note: L3 only mode blocks all layer 2 protocols, including ARP, when they end with a deny any statement.

Step	Action
To set the global ACL mode	
1	Set the ACL mode: `access-list set filter-mode <l2-l3-combo|l3-only>`
To set the global ACL mode to default	
2	Set the ACL mode back to default: `access-list unset {[filter-mode]}`

—end—

Example

This example sets the filter mode to l2-l3 combo:

```
access-list-list set filter-mode l2-l3 combo
```


Copyright© 2022 Ciena® Corporation

Procedure 3-2
Creating an ACL profile

You can

- create an ACL profile
- delete a single ACL profile, a list of ACL profiles or all ACL profiles.

  Note: An ACL profile can only be deleted if it has not been applied to any interfaces.

Step	Action

To create an ACL profile

1 Create an ACL profile:

```
access-list create acl-profile <String[1..31]> {default-filter-action <allow|deny>}
```

where

acl-profile <String[1..31]>	the name of the ACL profile being created.
default-filter-action <allow | deny>	sets the default filter action.

To delete a single ACL profile, a list of ACL profiles or all ACL profiles

2 Delete an ACL profile if it has not been applied to any other interfaces:

```
access-list delete <ACL Profile List>
```

where

profile <ACL Profile List>	is the single ACL profile, a list of ACL profiles or all profiles you are deleting. To specify all ACL profiles, specify “all”.

—end—

Example

This example creates the ACL profile ac1_1 with the filter action of deny.

```
access-list create acl-profile ac1_1 default-filter-action deny
```


Copyright© 2022 Ciena® Corporation

Procedure 3-3
Adding or removing ACL rules to or from an ACL profile

ACL rules are always considered to be inactive when the ACL profile you are adding the rules to is not attached to an interface.

You can remove a

- single rule
- list of rules
- all rules in the profile by specifying “all”, with the exception of the default

Note: Any filter term not explicitly specified is interpreted as fully wild-carded or all. This is also true for the is-fragment attribute, which if not specified allows the rule to match on both fragmented and unfragmented packets.

Step	Action

To add an ACL profile rule to the profile

1 Add an ACL profile to the profile:

```
access-list add profile <ACL Profile> rule <acl-rule-create-object> {precedence <NUMBER: 0..254>}, {filter-action <allow|deny>}, {{[src-mac <MAC address with mask>][dst-mac <MAC address with mask>], [base-etype <E-type Value in Hex>], [vtag-stack <vlan-stack-object>], [src-ip <IP address in CIDR notation>],[dst-ip <IP address in CIDR notation>], [dscp <DSCP with mask>], [ip-protocol <NUMBER: 0..255>], [is-fragment <yes|no>], [l4-src-port <L4 port range: Min-Max>],[l4-dst-port <L4 port range: Min - Max>], [l4-dst-protocol <l4-dst-protocol-attr>], [tcp-flags <tcp-flags-attr>]}| [any]}
```

where

profile <acl-profile-object-list>	is the ACL profile name.
rule <acl-rule-create-object>	is the mandatory ACL rule name. The ACL rule name is a string identifier to allow you to manipulate the rule.
precedence <NUMBER: 0..254>	is the mandatory ACL rule precedence, with 0 as the highest and 255 as the lowest. The precedence value of 255 is reserved for the default rule.

Copyright© 2022 Ciena® Corporation

where

filter-action <allow|deny>	specifies the mandatory filter action to perform when the rule is matched: allow or deny.
[src-mac <MAC address with mask>]	is the source MAC address with mask. This represents the source MAC address match criteria for the ACL filter. The MAC mask is optional. You can also use specific MACs. The mask for MAC addresses is the same format as the MAC addresses themselves. It is specified by adding a “/” following the MAC. For example: `access-list add profile <profile name> rule MACmask precedence 1 filter-action deny src-mac 00:03:18:be:02:b0/ff:ff:ff:ff:00:00`
[dst-mac <MAC address with mask>]	is the destination MAC address with mask. This represents the destination MAC address match criteria for the ACL filter. The MAC mask is optional. You can also use specific MACs. The mask for MAC addresses is the same format as the MAC addresses themselves. It is specified by adding a “/” following the MAC. For example: `access-list add profile <profile name> rule MACmask precedence 1 filter-action deny dst-mac 00:03:18:be:02:b0/ff:ff:ff:ff:00:00`
[base-etype <E-type Value in Hex>]	specifies the base ethertype value to match on.
[vtag-stack <vlan-stack-object>]	specifies the VID (off the wire), PCP and DEI values to match on. Format: <vid[/mask]>.<pcp[/mask]>.<dei>. A "*" value may be specified for the "vid" and "pcp" values to accept any vid and/or pcp value.
[src-ip <IP address in CIDR notation>]	specifies the source IPv4/IPv6 address (with mask) to match on.
[dst-ip <IP address in CIDR notation>]	specifies the destination IPv4/IPv6 address (with mask) to match on.
[dscp <DSCP with mask>]	species the DSCP value to match on.
[ip-protocol <NUMBER: 0..255>]	specifies the IP protocol to match on.
[is-fragment <yes|no>]	specifies the VID (off the wire), PCP and DEI values to match on. Format: <vid[/mask]>.<pcp[/mask]>.<dei>. A "*" value may be specified for the "vid" and "pcp" values to accept any vid and/or pcp value.

Copyright© 2022 Ciena® Corporation

where

[l4-src-port <L4 port range: Min-Max>]	is the L4 source port. This can be a single value or a dashed range (min - max).
[l4-dst-port <L4 port range: Min - Max>]	is the L4 destination port. This can be a single value or a dashed range (min - max).
[l4-dst-protocol <l4-dst-protocol-attr>]	is the L4 destination port which is mutually exclusive with the l4-dst-protocol.
[tcp-flags <tcp-flags-attr>]	are the tcp flags which are defined with one or more of the flags in a comma-separated list.
[any]	is mutually exclusive with all other filter terms.

To remove a single ACL rule, a list of ACL rules or all ACL rules from the profile

2 Remove a single ACL rule, a list of ACL rules, or all ACL rules from the profile:

```
access-list remove profile <ACL Profile> rule <acl-rule-name or "all">
```

where

profile <ACL Profile>	specifies the ACL profile the ACL rule is being removed from.
rule <acl-rule-name or "all">	specifies the ACL rule that is being removed from the profile, or to remove all ACL rules.

—end—

Example

This example adds a rule ID named rule-1 to allow telnet traffic from host IP 192.168.1.1:

```
access-list add acl-profile acl_1 rule rule-1 precedence 10 filter-action
allow src-ip 192.168.10.1/32 l4-dst-protocol telnet
```

This example then adds another lower priority rule to allow ARP packets from a specific source MAC address:

```
access-list add acl-profile acl_1 rule-id rule-2 precedence 20 filter-action
below src-mac 00:02:5a:01:d2:f8 etype 0x806
```

This example removes an ACL rule from an ACL profile.

```
access-list remove profile acl_1 rule rule-2
```


Copyright© 2022 Ciena® Corporation

Procedure 3-4
Applying ACL profiles

To make the ACL rules active, you must apply the ACL profile to a logical interface, such as:

- VLAN
- port
- virtual switch
- L3 interface

Step	Action

To apply an ingress ACL profile on the VLAN

1 Apply an ingress ACL profile on the VLAN:

```
vlan set vlan <vlan-object> [ingress-acl <acl-profile-object-list>]
```

where

vlan <vlan-object>	is the VLAN that you are applying the ingress ACL profile to.
ingress-acl <acl-profile-object-list>	is the ACL profile to be applied on the ingress direction of this VLAN.

To remove the ingress ACL profile from the VLAN

2 Remove the ingress ACL profile from the VLAN:

```
vlan unset vlan <vlan-object> ingress-acl
```

where

vlan <vlan-object>	is the VLAN that are removing from the ingress ACL profile.
ingress-acl	is the ingress ACL profile being removed from the ingress direction of this VLAN.

Copyright© 2022 Ciena® Corporation

To apply an ingress ACL profile on the port

3 Apply an ingress ACL profile on the port:

```
port set port <port> ingress-acl <acl-profile-object-list>
```

where

<port>	is the port that you are applying the ingress ACL profile to.
[ingress-acl <acl-profile-object-list>]	is the ingress ACL profile being applied to the port.

To remove the ingress ACL profile from the port

4 Remove the ingress ACL profile from the port:

```
port unset port <port> ingress-acl <acl-profile-object-list>
```

where

<port>	is the port that you are removing from the ingress ACL profile.
ingress-acl <acl-profile-object-list>]	is the ingress ACL profile being removed from the port.

To apply an ACL profile on a virtual switch

5 Apply an ACL profile on a virtual switch:

```
virtual-switch ethernet set vs <virtual switch> ingress-acl <ingress-acl>
```

where

vs <virtual switch>	is the name of the virtual switch.
ingress-acl <ingress-acl>	is the name of the ingress ACL.

To remove an ACL profile on a virtual switch

6 Remove an ACL profile on a virtual switch:

```
virtual-switch ethernet unset vs <virtual switch> ingress-acl <ingress-acl>
```


Copyright© 2022 Ciena® Corporation

To apply an ACL profile on an L3 interface

7 Apply an ACL profile on an L3 interface:

```
interface set ip-interface <ip-interface> ingress-acl
<acl-profile-object-list>
```

where

ip-interface <ip-interface>	is the name of the L3 interface.
ingress-acl <acl-profile-object-list>]	is the ingress ACL profile being applied to the L3 interface.

To remove an ACL profile from an L3 interface

8 Remove an ACL profile from an L3 interface:

```
interface unset ip-interface ingress-acl
```

—end—

Example

This example applies an ACL profile named acl_1 on the ingress traffic of port 1:

```
port set port 1 ingress-acl acl_1
```

This example applies the ACL profile named acl1 to the ingress traffic on an IP interface named ipint:

```
interface set ip-interface ipint ingress-acl acl1
```


Copyright© 2022 Ciena® Corporation

Procedure 3-5
Viewing ACL information

You can view

- global ACL feature information, including the administrative state and number of ACL profiles configured
- the maximum number of unique profile and rule definitions that can be configured on the device
- ACL profile and rule configuration information

Step	Action

To view global ACL feature information

1 View global ACL feature information:

```
access-list show
```

To view the maximum number of unique profile and rule definitions that can be configured on the device

2 Display the limit on the maximum number of unique profile and rule definitions that can be configured on the device:

```
access-list show [resources]
```

where

resources	displays the maximum number of unique profile and rule definitions that can be configured on the device. This attribute is independent of hardware resource usage.

To view ACL profile and rule configuration information

3 Display ACL profile and rule configuration information:

```
access-list show profile <profile> [statistics] [port
<Port list>] [vlan <VLAN ID List>] [vs <Virtual Switch
Ethernet Name>] [ip-interface <Interface Name>]
[interface <remote | local>] [all-instances] [active]
[brief] [interfaces]
```

where

profile <profile>	is the ACL profile that you are displaying rule statistics for.
statistics	retrieves the global hit statistics for each rule definition. The hit statistics displayed are the global hit statistics for all instances of this profile. You must additionally specify one interface or all-instances.
port <port List>	is the port or ports to show the profile instance statistics for.

Copyright© 2022 Ciena® Corporation

where

vlan <VLAN ID List>	is the VLAN or VLANs to show the profile instance statistics for.
vs <Virtual Switch Ethernet Name>	is the virtual switch.
ip-interface <Interface Name>	is the IP-interface to show the profile instance statistics for.
interface <remote | local>	specifies the local or remote interface.
all-instances	specifies all attached interfaces.
active	displays active rules only, that is, non-zero statistics.
brief	displays summarized output, that is, without filter terms.
interfaces	displays attached interfaces.

—end—

Example

This example shows sample output for the access-list show command.

```
>access-list show

+----------- ACL GLOBAL CONFIGURATION ------------+
| Parameter                | Value                  |
+--------------------------+------------------------+
| Admin State              | Enabled                |
| Oper State               | Enabled                |
| Filter Mode              | L2/L3 Combo            |
| Number of ACL Profiles   | 1                      |
+--------------------------+------------------------+

+---------------- ACL PROFILES ----------------+
| Profile | Admin | Oper  | Num   | Attached   |
| Name    | State | State | Rules | Interfaces |
+---------+-------+-------+-------+------------+
| prof1   | Ena   | Dis   | 2     | 0          |
+---------+-------+-------+-------+------------+

+---------------- ACL PROFILES ----------------+
| Profile | Admin | Oper  | Num   | Attached   |
| Name    | State | State | Rules | Interfaces |
+---------+-------+-------+-------+------------+
| ad      | Ena   | Dis   | 1     | 0          |
+---------+-------+-------+-------+------------+
```

This example shows sample output for the access-list show resources command:

Copyright© 2022 Ciena® Corporation

```
>access-list show resources

+------------------------- ACL RESOURCE USAGE --------------------------+
| Parameter                            | Value                            |
+--------------------------------------+----------------------------------+
| Number of ACL Profiles Defined       | 1                                |
| Remaining ACL Profile Definitions    | 1023                             |
| Number of ACL Rules Defined          | 2                                |
| Remaining ACL Rule Definitions       | 2046                             |
| Number of ACL Active Rules Used      | 0                                |
| Number of ACL Active Rules Free      | 3584                             |
+--------------------------------------+----------------------------------+
```

This example shows sample output for the profile ACL_1:

```
> access-list show profile ACL_1

+----------------------- ACL PROFILE INFO -------------------------+
| Parameter                        | Value                          |
+----------------------------------+--------------------------------+
| Name                             | ACL_1                          |
| Index                            | 1                              |
| Admin State                      | Enabled                        |
| Oper State                       | Enabled                        |
| Default Filter Action            | Deny                           |
| Number of Rules in Profile       | 2                              |
| Number of Attached Interfaces    | 8                              |
+----------------------------------+--------------------------------+

+------------------------------- ACL PROFILE RULES ----------------------------------+
| Rule     | Rule | Filter |        Filter Terms                                      |
| Name     | Prec | Action |   Type        |     Value                                  |
+----------+------+--------+---------------+-------------------------------------------+
| rule-1   | 10   | Allow  | Src-IP        | 192.168.1.1/24                            |
|          |      |        | IP-Proto      | 6                                         |
|          |      |        | DSCP          | 25                                        |
|          |      |        | L4-Src-Port   | 2052                                      |
+----------+------+--------+---------------+-------------------------------------------+
| default  | 255  | Deny   | Any           | -                                         |
+----------+------+--------+---------------+-------------------------------------------+
```

This example shows sample output for the interfaces on profile ACL_1:

```
> access-list show profile ACL_1 interfaces

+--------- ACL PROFILE INTERFACE ATTACHMENTS ----------+
| Interface Type    | Interface Name/ID     | Direction |
+-------------------+-----------------------+-----------+
| Port              | 1                     | Ingress   |
+-------------------+-----------------------+-----------+
| Port              | 2                     | Ingress   |
+-------------------+-----------------------+-----------+
| VLAN              | VLAN#100              | Ingress   |
+-------------------+-----------------------+-----------+
| VLAN              | VLAN#200              | Ingress   |
+-------------------+-----------------------+-----------+
| Virtual Switch    | VS_1                  | Ingress   |
+-------------------+-----------------------+-----------+
| IP Interface      | IP_INT_1              | Ingress   |
+-------------------+-----------------------+-----------+
| Local Interface   | -                     | Ingress   |
+-------------------+-----------------------+-----------+
| Remote Interface  | -                     | Ingress   |
+-------------------+-----------------------+-----------+
```

This example shows sample output for statistics on profile ACL_1:

```
> access-list show profile ACL_1 statistics
```


Copyright© 2022 Ciena® Corporation

```
+---------------------- ACL PROFILE GLOBAL RULE STATISTICS ----------------------------------+
| Rule | Filter |        Filter Terms                          |    Hit Statistics              |
| Prec | Action |   Type        |     Value                    |      Packets       |    Bytes  |
+------+--------+---------------+------------------------------+--------------------+-----------+
| Rule Name: rule-1                                                                             |
|-----------------------------------------------------------------------------------------------|
| 10   | Allow  | Src-IP        | 192.168.1.1/24               |               30 |        2160 |
|      |        | IP-Proto      | 6                            |                  |             |
|      |        | DSCP          | 25                           |                  |             |
|      |        | L4-Src-Port   | 2052                         |                  |             |
+------+--------+---------------+------------------------------+--------------------+-----------+
| Rule Name: default                                                                            |
|-----------------------------------------------------------------------------------------------|
| 255  | Deny   | Any           | -                            |               21 |        1512 |
+------+--------+---------------+------------------------------+--------------------+-----------+
```

This example shows sample output for brief statistics on profile ACL_1:

```
>access-list show profile ACL_1 statistics brief
```

```
+----------------- ACL PROFILE GLOBAL RULE STATISTICS ------------------+
| Rule     | Rule | Filter |                Hit Statistics               |
| Name     | Prec | Action |        Packets       |        Bytes         |
+----------+------+--------+----------------------+----------------------+
| rule-1   | 10   | Allow  |                   30 |                 2160 |
+----------+------+--------+----------------------+----------------------+
| default  | 255  | Deny   |                   21 |                 1512 |
+----------+------+--------+----------------------+----------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 3-6
Clearing ACL rule hit statistics

You can clear ACL rule hit statistics for all rules in an ACL profile. This command clears the hit statistics globally for

- all instances if you do not specify an interface, or
- specific profile instances if an interface is specified

Step	Action
1	Clear ACL rule hit statistics `access-list clear profile <acl-profile-object-list> {{statistics} [[port <port-object-list>]|[vlan <vlan-object-list>]|[vs <vs-object>]|[remote-interface]|{local-interface]},` where profile <acl-profile-object-list> — is the ACL profile you are clearing port <port-object-list> — is the port associated with the ACL profile. vlan <vlan-object-list> — is the VLAN associated with the ACL profile. vs <vs-object> — is the virtual switch associated with the ACL profile [remote-interface] | [local-interface]} — is the remote or local interface associated with the ACL profile. **—end—**

Example

This example clears ACL profile instance stats on port 1:

```
access-list clear profile acl_1 statistics port 1
```


Copyright© 2022 Ciena® Corporation

Procedure 3-7
Enabling and disabling an ACL profile

An ACL profile is enabled on creation.

Step	Action

To enable an ACL profile

1 Enable an ACL profile:

```
access-list enable profile <acl-profile-object-list>
```

where

profile <acl-profile-object-list>	is the ACL profile you are enabling.

To disable an ACL profile

2 Disable an ACL profile:

```
access-list disable profile <acl-profile-object-list>
```

where

profile <acl-profile-object-list>	is the ACL profile you are disabling

—end—

Copyright© 2022 Ciena® Corporation

Procedure 3-8
Enabling and disabling ACL

You can enable or disable ACL globally.

Step	Action
To enable ACL globally	
1	Enable an ACL globally: `access-list enable`
To disable ACL globally	
2	Disable ACL globally: `access-list disable`

—end—

Copyright© 2022 Ciena® Corporation

Hardware resource management

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

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

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

This customized resource management is supported for these features:

- Access control list
- Advanced Layer 3 features (IRB)
- Broadcast containment
- CFM
- Loss measurement
- MPLS class type to queue group mapping
- Traffic profiling
- Transport OAM
- Virtual circuit statistics
- Virtual circuit transforms (VC transforms)
- Virtual switch Layer 2 enhanced transforms

Copyright© 2022 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 on the platform as shown in this table.

Table 4-1
Resource pools

Platform	Resource type	Number of pools	Pool boundary
3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928	Classifiers	8	256
	Meters	1 (global)	2048
	Counters	8	256
3942	Classifiers	16	256
	Meters	1 (global)	4096
	Counters	1 (global)	4096
5142, 5160	Classifiers	16	512
	Meters	1 (global)	8192
	Counters	1 (global)	8192

Resources are allocated differently depending on the platform.

On the 3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928 platforms:

- Classifier and counter resources are reserved on a per-feature basis through configuration.
- Counters are reserved, but not configurable; configuring classifiers reserves the same number of counters.
- Meter resources are allocated on a first-come, first-served basis from a global pool and are not reservable through configuration.

On the 3942, 5142, and 5160 platforms:

- Classifier resources are reserved on a per-feature basis through configuration.
- Meter and counter resources are each allocated on a first-come, first-served basis from a global pool and are not reservable through configuration.

Copyright© 2022 Ciena® Corporation

Each feature consumes hardware resources at a different rate depending on the configuration as shown in this table.

Table 4-2
Resources consumed by feature and type

Feature	Resource type	Resources consumed
Access control list	Classifiers	2
	Meters	0
	Counters	1
Broadcast containment	Classifiers	Varies depending on mode (resource mode or non-resource mode) and device. See Scale management.
	Meters	1
	Counters	1
CFM (configurable on 5142 and 5160 only)	Classifiers	On 5142 and 5160 platforms, CFM consumes 61 static entries by default. For details, refer to *39XX/51XX Switches and Platforms Fault and Performance Management*.
	Meters	0
	Counters	0
Loss Measurement	Classifiers	6
	Meters	0
	Counters	0
MPLS queuing	Classifiers	The mpls-queuing feature is used for directing MPLS transit and (non-stitched) L2-VPN traffic to the egress queue group according to the class type profile attached to the tunnel's IP interface. This feature is available on 3924, 3926, 3928, 3942, 5142 and 5160 devices only. When enabled, each CT-mapped transit tunnel consumes one classifier; • Each non-stitched L2-VPN over CT-mapped unprotected tunnel consumes one classifier. • Each non-stitched L2-VPN over CT-mapped protected tunnel consumes two classifiers.
	Meters	0
	Counters	0

Copyright© 2022 Ciena® Corporation

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

Feature	Resource type	Resources consumed
MPLS MSPW queuing	Classifiers	Varies. This feature directs stitched L2-VPN traffic to egress queue group according to the class type profile attached to the tunnels IP interface. This feature is available on 3924, 3926, 3928, 3942, 5142 and 5160 devices. When enabled, each stitched L2-VPN over CT-mapped unprotected tunnel consumes one classifier; each stitched L2-VPN over CT-mapped protected tunnel consumes two classifiers.
	Meters	0
	Counters	0
Traffic profiling	Classifiers	Standard Traffic Profiles: 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 — Consumes classifiers based on the sum of the parent and child mode classifiers. • hierarchical-vlan — Consumes classifiers based on 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 are 12.
		Global Traffic Profiles: 2
	Meters	2
	Counters	2

Copyright© 2022 Ciena® Corporation

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

Feature	Resource type	Resources consumed
Transport OAM (not configurable on 3903, 3904, 3905, 3906 devices)	Classifiers	On 3924, 3926, 3928, 3942, 5142, 5160 devices: • 54 static entries, dynamic entries vary per feature
	Meters	0
	Counters	0
Virtual circuit statistics	Classifiers	Varies. When statistics are enabled for a virtual circuit, classifiers are consumed depending on the number of ports in the provider VLAN associated with the virtual circuit and the platform. A virtual circuit consumes one classifier for received traffic, and each port in the provider VLAN consumes one classifier for transmitted traffic.
	Meters	0
	Counters	Varies. When statistics are enabled for a virtual circuit, counters are consumed depending on the number of ports in the provider VLAN associated with the virtual circuit. A virtual circuit consumes one counter for received traffic, and each port in the provider VLAN consumes one counter for transmitted traffic.
Virtual circuit transforms	Classifiers	On 3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928, devices: • 160 static entries by default On 3942 devices: • 158 static entries by default On 5142 and 5160 devices: • 286 static entries by default Each VC assigned a VC transform consumes 1 dynamic entry.
	Meters	0
	Counters	0
Virtual switch enhanced L2 transform (not configurable on 3924, 3926, 3928, 3942, 5142, 5160 devices)	Classifiers	17 static entries by default
	Meters	0
	Counters	0

Copyright© 2022 Ciena® Corporation

Hardware management allows you to free up resources used by other features to be used elsewhere. The management pool must be configured correctly to share these resources.

Virtual port multicast group resources

Virtual Port Multicast (VPMC) Groups are shared by various features such as sub-port virtual switches, VPLS virtual switches, IGMP Multicast Groups. Each feature instance consumes one virtual port multicast group resource.

> ***Note:*** If a feature is not supported, a user is never allowed to configure an instance on it.

This table describes the virtual port multicast group resources.

Table 4-3
VPMC platform scale

Platform	Maximum Count	System Reserved	Usable
3924, 3926, 3928	256	6	250
3903, 3903x, 3906	512	3	509
3904, 3905	512	4	508
3942	8192	7	8185
5142, 5160	16384	12	16372

The system reserves these virtual port multicast groups at the time of device bootup:

- Two virtual port multicast groups are reserved by hardware for 3924/3926/3928 and one for other platforms.
- One virtual port multicast group is reserved to drop packets for IGMP services over QinQ and VPLS VS.
- One virtual multicast port group is reserved by platforms that support TDM.
- One virtual multicast port group is reserved for every 1000 MPLS-VC supported on the system for blocking pseudowire.

Copyright© 2022 Ciena® Corporation

Procedures

The procedures for hardware resources are:

- “Configuring resources” on page 4-8
- “Configuring hardware access control list resources” on page 4-10
- “Freeing all hardware access control list resources” on page 4-12
- “Configuring broadcast containment resources” on page 4-13
- “Freeing all broadcast containment resources” on page 4-15
- “Configuring CFM resources” on page 4-16
- “Freeing all CFM resources” on page 4-17
- “Configuring loss measurement resources” on page 4-18
- “Freeing all loss measurement resources” on page 4-20
- “Configuring traffic profiling resources” on page 4-21
- “Freeing all traffic profiling resources” on page 4-23
- “Configuring transport OAM resources” on page 4-24
- “Freeing all transport OAM resources” on page 4-26
- “Configuring virtual circuit statistics resources” on page 4-27
- “Freeing all virtual circuit statistics resources” on page 4-29
- “Configuring virtual circuit transform resources” on page 4-30
- “Freeing all virtual circuit transform resources” on page 4-32
- “Configuring virtual switch L2 enhanced transform resources” on page 4-33
- “Freeing all virtual switch L2 enhanced transform resources” on page 4-35
- “Displaying resource configuration information” on page 4-36
- “Addressing classifier resource allocation too small for current configuration error” on page 4-37
- “Displaying resource configuration in the configuration file” on page 4-39
- “Displaying system resources” on page 4-40

Copyright© 2022 Ciena® Corporation

Procedure 4-1
Configuring resources

Use this procedure to configure 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 reserved on the 5160 platform, 750 is rounded up to the nearest pool boundary, 1024, where two classifier resource pools are allocated and reserved.

Three types of resources can be reserved or allocated:

Table 4-4
Resource types for reservation or allocation

Resource types	Reservation or allocation rules
Classifier	Classifier resources must be reserved through configuration on all SAOS 6.x platforms.
Meter	Meter resources are allocated from a global pool on all SAOS 6.x platforms on a first-come, first-served basis and cannot be reserved for a feature.
Counter	Counter resource allocation varies based on the platform. On 3903, 3903x, 3904, 3905, 3906, 3616, 3924, 3926, 3928 platforms, counter resources are required reserved, and are reserved when classifier resources are configured. On 3942, 5142, and 5160 platforms, counter resources are allocated from a global pool on a first-come, first-served basis and cannot be reserved for a feature.

The resource manager provides the current usage of each resource type for each feature regardless if the resource is reserved for the feature or taken from a global pool of 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.

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

Copyright© 2022 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 resource pool count for each feature you want to reassign resources from. `resource-manager pool set resource <classifier> feature <access-list | advanced-l3 | broadcast-containment | cfm | loss-measurement | mpls-queuing | mpls-mspw-queuing | traffic-profiling | transport-oam | vc-statistics | vc-transform | vs-enhanced-l2-transform> count <NUMBER>`
2	Set the resource pool count for each feature you want to reassign resources to. `resource-manager pool set resource <classifier> feature <access-list | advanced-l3 | broadcast-containment | cfm | loss-measurement | mpls-queuing | mpls-mspw-queuing | traffic-profiling | transport-oam | vc-statistics | vc-transform | vs-enhanced-l2-transform> count <NUMBER>`
3	Validate the configuration. `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
4	Save the configuration. `configuration save`
5	Reboot the system to implement the changes. `chassis reboot`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-2
Configuring hardware access control list resources

By default, the hardware access control list (ACL) feature does not have any assigned resources. To use this feature, you need to assign available classifier and counter resources to it.

You can also use this procedure to restore hardware ACL resources to default values if you allocated the resources to different features, and then choose to use hardware ACL.

The maximum number of resources that can be reserved for ACLs are shown in this table.

Table 4-5
ACL default and maximum resource reservation by platform

Platform	Default	Maximum resources
3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928	0	1536
3942	0	3328
5142, 5160	0	7168

Step	Action
1	Set the pool count for classifier resources to use the ACL feature on the platform. `resource-manager pool set resource classifier feature access-list count <ClassifierResources>` where <ClassifierResources> is the pool count for classifier resources for the platform. ***Note:*** To restore hardware ACL resources to default values, enter the default pool count for classifier resources for the platform.
2	Validate the configuration. `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
3	Save the configuration. `configuration save`

Copyright© 2022 Ciena® Corporation

4 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-3
Freeing all hardware access control list resources

To free up hardware access control list (ACL) resources, all ACL profiles must be detached from all interfaces that are using them. It is recommended that the ACL profiles also be deleted, but it is not required to free up classifier resources.

Step	Action
1	Detach all ACL profiles from the ports that are using them: `port unset port <port-number> ingress-acl`

2	Detach all ACL profiles from all VLANs/virtual switches that are using them:
	`vlan unset vlan <vlan-id> ingress-acl`
	`virtual-switch ethernet unset vs <vs-name> ingress-acl`

3	Detach all ACL profiles from all L3 interfaces that are using them:
	`interface unset ip-interface <intf-name> ingress-acl`
	`interface remote unset ingress-acl`
	`interface local unset ingress-acl`

4	Delete all ACL profiles on the device (optional): `access-list-delete profile all`
5	Set the pool count for the classifier resources to 0: `resource-manager pool set resource classifier feature access-list count 0`
6	Validate the configuration: `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
7	Save the configuration: `configuration save`
8	Reboot the system to implement the changes: `chassis reboot`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-4
Configuring broadcast containment resources

Use this procedure to configure broadcast containment resources.

You can also use this procedure to restore broadcast containment resources to default values if you allocated the resources to different features, and then choose to use broadcast containment.

This table shows the default and maximum number of resources that can be reserved for broadcast containment.

Table 4-6
Broadcast containment default and maximum resources reservation by platform

Platform	Default	Maximum resources
3903, 3903x, 3904, 3905, 3906	256	256
3924, 3926, 3928	0	1,536
3942	256	3,328
5142, 5160	512	7,168

Step	Action
1	Set the pool count for classifier resources to use the broadcast containment feature on the platform: `resource-manager pool set resource classifier feature broadcast-containment count <ClassifierResources>` where <ClassifierResources> is the desired pool count for classifier resources for the platform. ***Note:*** To restore broadcast containment resources to default values, enter the default pool count for classifier resources for the platform.
2	Validate the configuration. `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.

Copyright© 2022 Ciena® Corporation

3 Save the configuration.

```
configuration save
```

4 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-5
Freeing all broadcast containment resources

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	Validate the configuration: `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
4	Save the configuration: `configuration save`
5	Reboot the system to implement the changes: `chassis reboot`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-6
Configuring CFM resources

CFM resources are only configurable on 5142 and 5160 platforms.

This table shows the default and maximum number of resources that can be reserved for CFM.

Table 4-7
CFM default and maximum resource reservation by platform

Platform	Default	Maximum resources
5142 5160	512	512

Step	Action
1	Set the pool count for classifier resources to use the CFM feature on 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.
2	Validate the configuration. `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
3	Save the configuration. `configuration save`
4	Reboot the system to implement the changes. `chassis reboot` **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 4-7
Freeing all CFM resources

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

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	Validate the configuration: `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
4	Save the configuration: `configuration save`
5	Reboot the system to implement the changes: `chassis reboot`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-8
Configuring loss measurement resources

Loss measurement resources are required to be configured for utilizing hardware-assisted Y.1731 loss measurement sessions. Y.1731 loss measurement sessions are supported on all platforms.

By default, loss measurement does not have any assigned resources. To use this feature, you need to assign available classifier and counter resources to it. On 3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928, and 3942 platforms, allocation of one resource block (of 256 classifiers) supports up to 42 loss measurement sessions. A maximum of 120 loss measurement sessions can be configured for the platform. On 5142 and 5160 platforms, allocation of one resource block (of 512 classifiers) supports up to 85 loss measurement sessions. A maximum of 255 loss measurement sessions can be configured for the platform.

You can also use this procedure to restore loss measurement resources to default values if you allocated the resources to different features, and then choose to use loss measurement.

The default and maximum number of resources that can be reserved for loss measurement are shown in this table.

Table 4-8
Loss measurement default and maximum resource reservations by platform

Platform	Default	Maximum resources
3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928, 3942	0	768
5142, 5160	0	1536

Copyright© 2022 Ciena® Corporation

Step	Action
1	Set the pool count for classifier resources for the loss-measurement feature on the platform. `resource-manager pool set resource classifier feature loss-measurement count <ClassifierResources>` where <ClassifierResources> is the desired pool count for classifier resources for the platform. ***Note:*** To restore loss measurement resources to default values, enter the default pool count for classifier resources for the platform.
2	Validate the configuration. `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
3	Save the configuration. `configuration save`
4	Reboot the system to implement the changes. `chassis reboot` **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 4-9
Freeing all loss measurement resources

If you are finished using loss measurement resources for utilizing hardware-assisted Y.1731 loss measurement sessions, you can reallocate the resources back to 0 and use them for another feature.

Note that 0 is the default value for resources.

Step	Action
1	Set the pool count for classifier resources to 0. `resource-manager pool set resource classifier feature loss-measurement count 0`
2	Validate the configuration: `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
3	Save the configuration: `configuration save`
4	Reboot the system to implement the changes: `chassis reboot`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-10
Configuring traffic profiling resources

Traffic profiling resources must be reserved in the resource manager to configure and operate standard traffic profiles and global traffic profiles.

Use this procedure to configure traffic profiling resources.

You can also use this procedure to restore traffic profiling resources to default values if you allocated the resources to different features, and then choose to use traffic profiling.

> ***Note:*** When configuring traffic-profiling resources, changes to resource counts are immediately activated and applied to the system, except when resource counts are changed from a non-zero value to a non-zero value, in which case, a configuration save and device reboot is required.

The default and maximum number of resources that can be reserved for traffic profiling by resource type are shown in this table.

Table 4-9
Traffic profiling default and maximum resource reservation

Platform	Default	Maximum resources
3903, 3903x, 3904, 3905, 3906	1024	1536
3924, 3926, 3928	0	1536
3942	2048	3329
5142, 5160	4096	7168

Copyright© 2022 Ciena® Corporation

Step	Action
1	Set the pool count for classifier resources for the traffic-profiling feature on the platform: `resource-manager pool set resource classifier feature traffic-profiling count <NUMBER>` where count <NUMBER> — is the resource count. ***Note:*** To restore traffic profiling resources to default values, enter the default pool count for classifier resources for the platform.
2	Validate the configuration: `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
3	Save the configuration: `configuration save`
4	Reboot the system to implement the changes: `chassis reboot`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-11
Freeing all traffic profiling resources

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

Note: Changing traffic-profiling resources to zero is immediately activated and applied to the system, and therefore does not require the configuration to be saved and the device to be rebooted.

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` **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 4-12
Configuring transport OAM resources

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

On the 3924, 3926 and 3928 platforms, transport OAM does not have any assigned resources by default. To use this feature, you need to assign available classifier resources to it.

You can also use this procedure to restore transport OAM resources to default values if you allocated the resources to different features, and then choose to use transport OAM.

The default and maximum number of resources that can be reserved for transport OAM are shown in this table.

Table 4-10
Transport OAM default and maximum resource reservations by platform

Platform	Default	Maximum resources
3942	256	512
5142, 5160	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. ***Note:*** To restore transport OAM resources to default values, enter the default pool count for classifier resources for the platform.
2	Validate the configuration. `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
3	Save the configuration. `configuration save`

Copyright© 2022 Ciena® Corporation

4 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-13
Freeing all transport OAM resources

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

Step	Action
1	Set the pool count for classifier resources to 0. `resource-manager pool set resource classifier feature transport-oam count 0`
2	Validate the configuration: `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
3	Save the configuration: `configuration save`
4	Reboot the system to implement the changes: `chassis reboot`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-14
Configuring virtual circuit statistics resources

Use this procedure to configure virtual circuit statistics resources.

You can also use this procedure to 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 this table.

Table 4-11
Virtual circuit statistics default and maximum resource reservation by platform

Platform	Default	Maximum resources
3903, 3903x, 3904, 3905, 3906	256	1536
3924, 3926, 3928	0	1536
3942	256	3328
5142, 5160	512	7168

Step	Action
1	Set the pool count for classifier resources for the platform. `resource-manager pool set resource classifier feature vc-statistics count <ClassifierResources>` where <ClassifierResources> is the pool count for classifier resources for the platform. ***Note:*** To restore virtual circuit statistics resources to default values, enter the default pool count for classifier resources for the platform.
2	Validate the configuration. `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
3	Save the configuration. `configuration save`

Copyright© 2022 Ciena® Corporation

4 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-15
Freeing all virtual circuit statistics resources

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

Step	Action
1	Disable statistics for each virtual circuit with statistics collection turned on. `virtual-circuit ethernet set vc <vc> statistics off` where vc <vc> 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	Validate the configuration: `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
4	Save the configuration: `configuration save`
5	Reboot the system to implement the changes: `chassis reboot`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-16
Configuring virtual circuit transform resources

Use this procedure to configure virtual circuit transform resources.

You can also use this procedure to restore virtual circuit transform resources to default values if you allocated the resources to different features, and then choose to use virtual circuit transform.

By default, virtual circuit transform does not have any assigned resources. To use this feature, you need to assign available classifier resources to it.

The default and maximum number of classifier resources that can be reserved for VC transform are shown in this table.

Table 4-12
VC L2 transform default and maximum resource reservations by platform

Platform	Default	Maximum resources
3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928, 3942	0	768
5142, 5160	0	1536

Step	Action
1	Set the pool count for classifier resources. `resource-manager pool set resource classifier feature vc-transform count <NUMBER>` where count <NUMBER> is the number of resources to assign. ***Note:*** To restore virtual circuit transform resources to default values, enter the default pool count for classifier resources for the platform.
2	Validate the configuration. `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
3	Save the configuration. `configuration save`

Copyright© 2022 Ciena® Corporation

4 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-17
Freeing all virtual circuit transform resources

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

Note that 0 is the default value for classifier resources.

Step	Action
1	If any VC transforms have been created, delete all VC transforms on the device. `virtual-circuit transform delete l2-transform <vc-transform-name>` where l2-transform <vc-transform-name> is the virtual circuit transform name.
2	Set the pool count for classifier resources to 0. `resource-manager pool set resource classifier feature vc-transform count <NUMBER>`
3	Validate the configuration: `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
4	Save the configuration: `configuration save`
5	Reboot the system to implement the changes: `chassis reboot`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-18
Configuring virtual switch L2 enhanced transform resources

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

You can also use this procedure to restore virtual switch L2 enhanced transform resources to default values if you allocated the resources to different features, and then choose to use virtual switch L2 enhanced transform.

Virtual switch L2 enhanced transform resources are not applicable to 3924, 3926, 3928, 3942, 5142, and 5160 devices.

The default and maximum number of resources that can be reserved for virtual switch L2 enhanced transform are shown in this table.

Table 4-13
Virtual switch L2 enhanced transform default and maximum resource reservations by platform

Platform	Default	Maximum resources
3903, 3903x, 3904, 3905, 3906	0	256

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. ***Note:*** To restore virtual switch L2 enhanced transform resources to default values, enter the default pool count for classifier resources for the platform.
2	Validate the configuration. `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.

Copyright© 2022 Ciena® Corporation

3 Save the configuration.

```
configuration save
```

4 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-19
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 resources back to 0 and use them for another feature.

Note that 0 is the default value for classifier 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	Validate the configuration: `resource-manager validate` ***Note:*** This step is optional as it is automatically performed when you save the configuration.
5	Save the configuration: `configuration save`
6	Reboot the system to implement the changes: `chassis reboot`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 4-20
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 becomes on successful configuration save and reboot).

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 traffic profiling feature	Perform step 5.

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 resource-manager feature:

```
resource-manager show feature { access-list | traffic-
profiling | vc-transform }
```

—end—

Copyright© 2022 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, this error is displayed when the resource-manager validate command is executed:

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

Example

This 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=256 used=6
ERROR: vc-statistics: counter resource allocation too small for current
configuration
ERROR:               : candidate=0 active=256 used=6
ERROR: Resource validation failed: Resource allocation invalid
```

This error condition is addressed by performing one of these actions:

- Increasing the pool count so that the number of resources matches or is greater than the current used count in 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 resources to the match the active value shown in the error message.

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

where

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

2 Validate the configuration.

```
resource-manager validate
```

Note: This step is optional as it is automatically performed when you save the configuration.

Copyright© 2022 Ciena® Corporation

3 Save the configuration.

```
configuration save
```

4 Reboot the system to implement the changes.

```
chassis reboot
```

To modify the configuration of the feature

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     | L2 Transform                     |
+------------------+---------------+-----------+----------------------------------+
| vc100            | 100           | On        | pvid1000                         |
| vc101            | 101           | On        | pvid1000                         |
| vc200            | 200           | On        |                                  |
+------------------+---------------+-----------+----------------------------------+
```

2 Modify the configuration of the feature.

```
virtual-circuit ethernet set vc <vc-name> statistics off
```

3 Validate the configuration.

```
resource-manager validate
```

Note: This step is optional as it is automatically performed when you save the configuration.

4 Save the configuration.

```
configuration save
```

5 Reboot the system to implement the changes.

```
chassis reboot
```

—end—

Copyright© 2022 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	View resource configuration in the configuration file: `configuration show` **—end—**

Example

This example shows sample output from the configuration show command.

```
configuration show

! 3942 Configuration File
!
!. . .
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
! RESOURCE CONFIG
!
resource-manager pool set feature broadcast-containment resource classifier count 0
resource-manager pool set feature vc-statistics resource classifier count 0
resource-manager pool set feature vs-enhanced-l2-transform resource classifier count 256
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
[more 9%] (q,g,space,enter)
```


Copyright© 2022 Ciena® Corporation

Procedure 4-23
Displaying system resources

You can display

- all system resource usage
- feature-wise usage of Virtual Port Multicast Groups
- 392X service-meter
- 392X service-counter

Step	Action

To display all system resources

1 Display all system resources:

```
resource-manager system-resources show
```

To display feature-wise usage of Virtual Port Multicast Groups

2 Display feature-wise usage of Virtual Port Multicast Groups:

```
resource-manager system-resources show virtual-port-
multicast-group
```

—end—

Example

This example shows sample output from the 392X resource-manager system-resources show command.

```
> resource-manager system-resources show
+------------------------ SYSTEM RESOURCES INFO -------------------------+
| Resource                      | Total | System Reserved | Used   | Free   |
+-------------------------------+-------+-----------------+--------+--------+
| Virtual Port Multicast Group  |   256 |               5 |    251 |      0 |
| Service Meter                 |  4096 |              28 |      4 |   4064 |
| Service Counter               |  2048 |              20 |      8 |   2020 |
+-------------------------------+-------+-----------------+--------+--------+
```

This example shows sample output from the 51XX resource-manager system-resources show command.

```
> resource-manager system-resources show

+------------------------ SYSTEM RESOURCES INFO -------------------------+
| Resource                      | Total | System Reserved | Used   | Free   |
+-------------------------------+-------+-----------------+--------+--------+
| Virtual Port Multicast Group  |   256 |               5 |    251 |      0 |
+-------------------------------+-------+-----------------+--------+--------+
```

This example shows sample output from the resource-manager system-resources show virtual-port-multicast-group command.

Copyright© 2022 Ciena® Corporation

```
> resource-manager system-resources show virtual-port-multicast-group

+-------- RESOURCE INFO --------+
| Parameter              | Value |
+------------------------+-------+
| Total                  |   256 |
| System Reserved        |     5 |
| Used                   |   251 |
| Free                   |     0 |
+------------------------+-------+
+------- RESOURCE USAGE --------+
| Feature                | Count |
+------------------------+-------+
| VPLS VS                |   251 |
| Sub-port VS            |     0 |
| IGMP Multicast Groups  |     0 |
+------------------------+-------+
```

This example shows sample output from the resource-manager system-resources show service-meter command.

```
> resource-manager system-resources show service-meter
+----- RESOURCE INFO -----+
| Parameter       | Value |
+-----------------+-------+
| Total           |  4096 |
| System Reserved |    28 |
| Used            |     4 |
| Free            |  4064 |
+-----------------+-------+
+--------- RESOURCE USAGE ---------+
| Feature         | Count | Max    |
+-----------------+-------+--------+
| Envelope Meter  |     0 |   4068 |
| Traffic Profile |     4 |    792 |
+-----------------+-------+--------+
```

This example shows sample output from the resource-manager system-resources show service-counter command.

```
*> resource-manager system-resources show service-counter
+----- RESOURCE INFO -----+
| Parameter       | Value |
+-----------------+-------+
| Total           |  2048 |
| System Reserved |    20 |
| Used            |     8 |
| Free            |  2020 |
+-----------------+-------+
+--------- RESOURCE USAGE ---------+
| Feature         | Count | Max    |
+-----------------+-------+--------+
| Envelope Meter  |     0 |   2028 |
| Traffic Profile |     8 |   1584 |
| MPLS            |     0 |   2028 |
+-----------------+-------+--------+
```


Copyright© 2022 Ciena® Corporation

Copyright© 2022 Ciena® Corporation

System timing configuration

System timing is the recovery and distribution of frequency, phase and time-of-day information to maintain synchronization between network elements.

System timing is performed by means of:

- Synchronous Ethernet
- External timing
  - BITS
  - GPS
- TDM line timing
- IEEE 1588 version 2 Precision Time Protocol
- Network Time Protocol (NTP) — For details on NTP, see *39XX/51XX Switches and Platforms Administration and Security.*

  Note: To configure Synchronous Ethernet (SyncE), external timing interfaces or TDM line timing, you need to install the Advanced Ethernet license key. To configure IEEE 1588 v2 Precision Time Protocol, you need to install the Advanced Synchronization license key. To obtain the Advanced Ethernet license key or the Advanced Synchronization license key, contact Ciena Sales.

System timing is supported on these platforms:

- 3904
- 3905
- 3924
- 3926
- 3928
- 5142
- 5160

  Note: The 3905, 3924, 3926, 3928 platforms do not support external timing interfaces.

Copyright© 2022 Ciena® Corporation

System timing recovers timing from external sources, that is, timing inputs, and distributes timing to external destinations, that is, timing outputs, where timing comprises frequency, phase and time-of-day.

A system can have multiple timing inputs: each timing input can have different characteristics that give it preference over other timing inputs. Different types of timing inputs can provide different components of frequency, phase and time-of-day to which the local clock is synchronized. In turn, the frequency, phase and time-of-day of the local clock can be distributed to other network elements through different types of timing outputs.

This table summarizes system timing inputs and outputs.

Table 5-1
System timing inputs and outputs

	Frequency	Phase	Time of Day
Inputs	• SyncE on any Ethernet port • IEEE 1588v2 on any Ethernet port • SYNC interface (BITS) • 10MHz interface • TDM port • Local oscillator	• IEEE 1588v2 on any Ethernet port • 1 PPS interface • SYNC interface	• IEEE 1588v2 on any Ethernet port • 1 PPS interface • SYNC interface • NTP • Set time
Outputs	• SyncE on any Ethernet port • IEEE 1588v2 on any Ethernet port • SYNC interface (BITS) • 10 MHz interface	• 1 PPS interface • SYNC interface • IEEE 1588v2 on any Ethernet port	• IEEE 1588v2 on any Ethernet port • 1 PPS interface • SYNC interface

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.

Copyright© 2022 Ciena® Corporation

There can be a number of Ethernet switches involved in the distribution of the reference timing signal. In such cases, the synchronization function in 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.

The Example of a synchronization network over synchronous Ethernet figure 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).

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

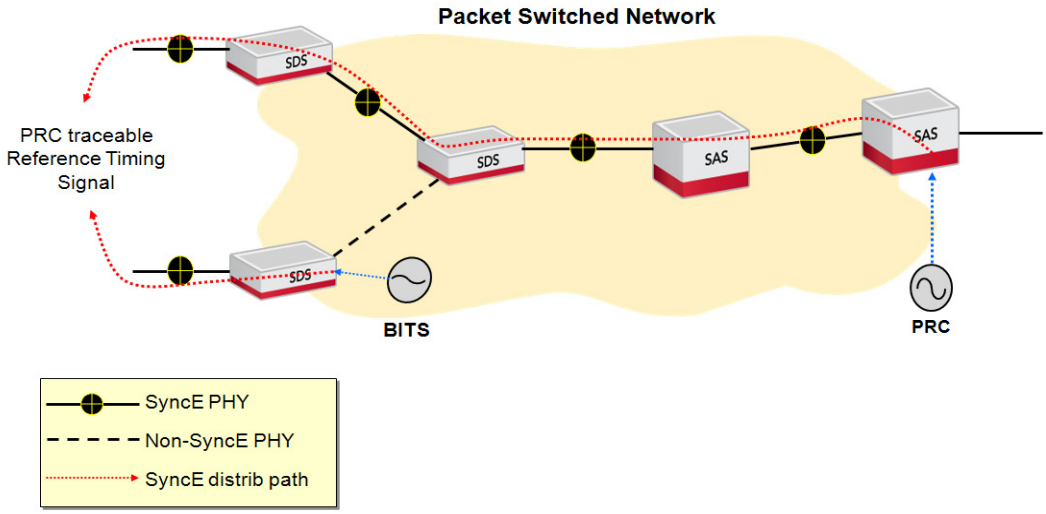

Synchronization Status Messaging (SSM) is used to convey priority and clock traceability. SSM messages indicate the quality level of the system clocks located in the various network elements. Quality level is the holdover performance of a clock.

SyncE port status is signaled by means of the Link Layer Discovery Protocol (LLDP). For more information on LLDP, see *39XX/51XX Switches and Platforms Base Configuration*.

Copyright© 2022 Ciena® Corporation

IEEE 1588 version 2 Precision Time Protocol

IEEE 1588v2 Precision Time Protocol (PTP) synchronizes real-time clocks over a network with sub-microsecond accuracy, very low network bandwidth, and minimal CPU usage. PTP supports synchronization over various types of network encapsulations, for example, UDP/IP and Ethernet. PTP supports multicast and unicast messaging.

Clocks are organized in a hierarchy: master clocks supply timing information to slave clocks. At the top of the hierarchy is the grandmaster clock. Typically, the grandmaster clock has a high-quality time source connected to it, for example, a GPS receiver or atomic clock.

39XX/51XX platforms can function as an ordinary clock (OC) slave as well as boundary clock (BC). On an OC, a single port participates in the clock hierarchy as a slave or master clock. On a BC, ports participate in the clock hierarchy as slave and master clocks.

In BC mode, on the master side of some 39XX/51XX platforms (currently 5142 and 5160 only), one-step timestamping is supported. One-step timestamping means that the outgoing PTP “sync” messages are timestamped by the hardware as the packet is being transmitted onto the wire. For two-step timestamping, the “sync” packet contains an estimate that is typically generated in the software. Two-step mode also sends a follow-up message with the exact time the outgoing PTP “sync” message left the device.

G.8275.1 profile support

This Ciena-only implementation of G.8275.1 includes the following elements:

- G.8275.1 protocol stack
- Clock recovery mechanism
- Use of the Alternate BMCA
- PTP interaction with EPL

G.8275.1 protocol stack

The G.8275.1 protocol stack runs directly in SAOS and makes use of a multi-thread mechanism to handle the protocol stack duties. One slave session per port is supported.

G.8275.1 has been supported on the 5142 and 5160 platforms since SAOS 6.17.1. SAOS platforms only support G.8275.1 T-BC (for phase/time of day) and G.8265.1 OC/BC for frequency. There is no support for 75.1 T-GM or 75.2 profile on any platforms.

Copyright© 2022 Ciena® Corporation

Clock recovery mechanism

The clock recovery mechanism only supports the G.8275.1 Telecom profile. It takes a timestamp 4-tuple (T1, T2, T3, T4) from the PTP stack and outputs a phase correction value once per second. For G.8275.1, it is assumed that SyncE is used for frequency synchronization.

This figure shows how the clock recovery mechanism uses the timestamp 4-tuple to calculate the offset and delay between the slave and master clocks.

Figure 5-2
Offset and delay calculations

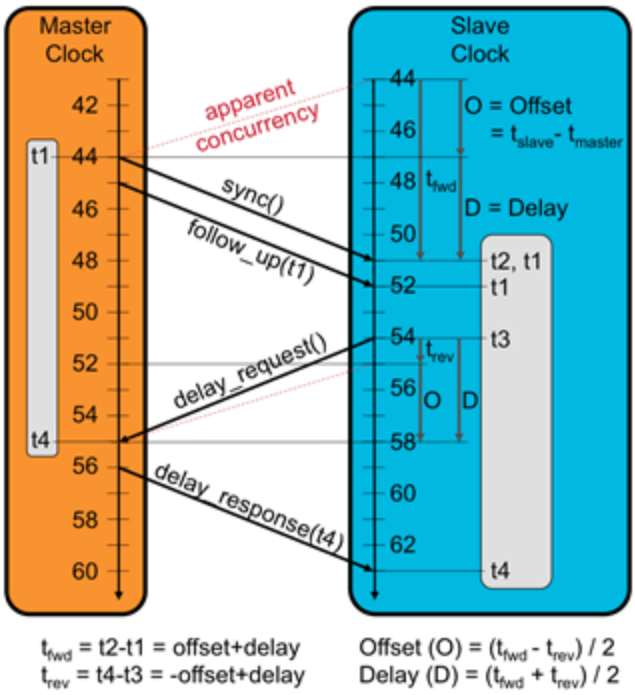

The offset value is used as the phase correction value output from the clock recovery mechanism.

Copyright© 2022 Ciena® Corporation

The clock recovery mechanism implements these states.

Table 5-2
Clock recovery mechanism states

State	Description
uninitialized	The initial default state on power-up.
init	The clock recovery mechanism is initializing.
learning	A transitory state in G.8275.1; sets initial jump detection limit.
acquiring	The clock recovery mechanism remains in this state until lock is achieved for the first time.
normal	“Normal” operating state. Locked to input phase reference.
holdover	Indicates loss of lock or expiry of “timestamp timeout” timer.

This figure details the possible state transitions.

Figure 5-3
Clock recovery mechanism state transitions

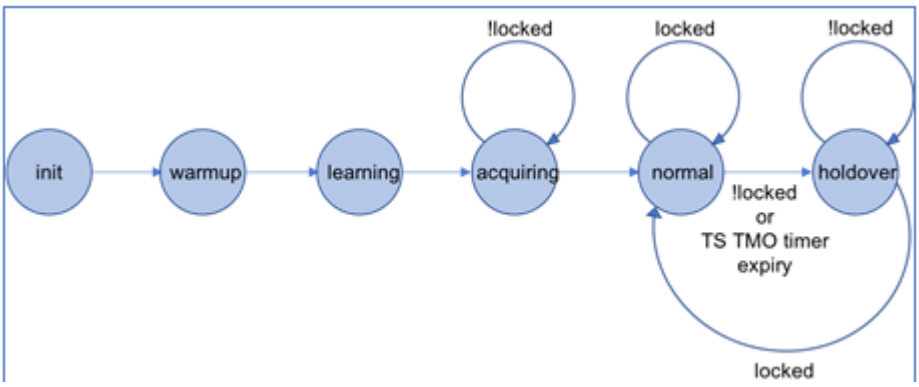

The clock recovery mechanism is completely driven by the receipt of timestamp 4-tuples. Phase adjustments and lock determination are only calculated when the clock recovery mechanism is presented with a timestamp 4-tuple. Since the clock recovery mechanism requires timestamp tuples to determine phase correction, an absence of timestamp tuples results in it not producing a phase correction value—effectively putting the clock recovery mechanism in holdover. The clock recovery mechanism maintains a one-second ‘timestamp timeout’ timer that is reset on receipt of a timestamp 4-tuple. If no timestamps are received within one second (a loss of 16 timestamp tuples with G.8275.1), the clock recovery mechanism enters holdover state. It remains in holdover, not producing any output, until timestamp tuples are again received and the clock recovery mechanism is locked to the incoming PTP phase clock.

Copyright© 2022 Ciena® Corporation

To determine lock, the clock recovery mechanism uses a running time deviation (TDEV) estimation with 1-, 10-, and 100-second observation intervals. Due to the nature of the TDEV calculation, three observation intervals are required before the TDEV value for a particular observation interval becomes valid. Since the clock recovery mechanism uses a 100-second TDEV observation interval, it takes a minimum of 300 seconds (five minutes) to initially determine lock state. Once a minimum of three observation intervals have passed, the lock state is updated on receipt of every timestamp 4-tuple.

Use of the Alternate BMCA

G.8275.1 specifies the use of the Alternate Best Master Clock Algorithm (BMCA), as originally described in clause 9.3.1 of the IEEE 1588-2008 specification.

For G.8275.1, it is required that physical input references are used for frequency and PTP packet input is used for phase. As such, the CLI enforces these restrictions:

- PTP input references cannot be added to frequency protection groups
- Only PTP input references can be added to phase protection groups

The Alternate BMCA (A-BMCA) is used to select from among the PTP input references. Its core is the data set comparison algorithm outlined in G.8275.1. With this algorithm, one phase clock is compared with another using the data sets representing these clocks. The best phase clock is then selected for synchronization.

Copyright© 2022 Ciena® Corporation

The clock class and frequency/time traceable flags output by Ciena's boundary clock are governed by Table V.3 from ITU-T G.8275.1.

Table 5-3
Applicable clockClass values

Phase/time traceability description	defaultDS. clockQuality. clockClass	frequencyTraceable flag	timeTraceable flag
When the boundary clock (T-BC) locks to the grandmaster clock (T-GM) and the T-GM is connected to a Primary Reference Time Clock (PRTC) in locked mode. For example, the PRTC is traceable to a Global Navigation Satellite System (GNSS). The T-BC transfers the T-GM clockclass and flags out.	6	TRUE	TRUE
When the T-BC locks to the T-GM, while the T-GM is in holdover, within the holdover specification, non-traceable to a Category 1 frequency source (see Note 1).	7	TRUE	TRUE
When the T-BC locks to the T-GM, while the T-GM is in holdover, within the holdover specification, non-traceable to a Category 1 frequency source (see Note 1).	7	FALSE	TRUE
When the T-BC is in holdover, within the holdover specification, traceable to a Category 1 frequency source (see Note 1).	135	TRUE	TRUE
When the T-BC is in holdover, within the holdover specification, non-traceable to a Category 1 frequency source (see Note 1).	135	FALSE	TRUE
When the T-BC locks to the T-GM, while the T-GM is in holdover, out of the holdover specification, traceable to a Category 1 frequency source (see Note 1).	140	TRUE	FALSE
When the T-BC locks to the T-GM, while the T-GM is in holdover, out of the holdover specification, traceable to a Category 2 frequency source (see Note 1).	150	FALSE	FALSE
When the T-BC locks to the T-GM, while the T-GM is in holdover, out of the holdover specification, traceable to a Category 3 frequency source (see Note 1).	160	FALSE	FALSE

Copyright© 2022 Ciena® Corporation

Table 5-3
Applicable clockClass values

Phase/time traceability description	defaultDS. clockQuality. clockClass	frequencyTraceable flag	timeTraceable flag
When the T-BC is in holdover, out of the holdover specification (see Note 1).	165	(see Note 2)	FALSE
When the T-BC has been without a time reference since start-up.	248	(see Note 2)	FALSE

Note 1: The holdover specification threshold controlling the time spent advertising clockClass values 7 or 135 could be set to zero so that the grandmaster or boundary clock would advertise a degraded clockClass value directly after losing traceability to a PRTC. In this case, initially after advertising clockClass values 140, 150, 160, or 165, a clock may still be within the holdover specification.

Note 2: The frequencyTraceable flag may be TRUE or FALSE, depending on the availability of a PRC-traceable physical layer frequency input signal.

Note 3: As an option, the clockClass range of a boundary clock can be extended from (135,165, 248) to (135, 140, 150, 160, 165, 248) for some cases, where (a) 140, 150, 160, and165 are related to the quality of the frequency reference, (b) the applicable circumstances of 140, 150, and 160 are the same as for the grandmaster clock, and (c) 165 corresponds to synchronous Ethernet equipment clock (EEC). If this option is used, then in a single PTP domain, all PTP clocks should implement this option (and should not be intermixed with clocks that do not implement this option).

The Category 1, 2, or 3 designations of Ciena's frequency sources are determined by this table.

Table 5-4
Mapping of clock QLs to Category 1, 2, and 3 frequency sources

Category (from Table 5-3)	ITU-T G.781 option I QLs	ITU-T G.781 option II QLs
Category 1 frequency source	QL-PRC	QL-PRS
Category 2 frequency source	QL-SSU-A	QL-ST2
Category 3 frequency source	QL-SSU-B	QL-ST3E

Ciena currently uses a five-minute threshold to determine whether or not a grandmaster or boundary clock is within holdover specification.

Copyright© 2022 Ciena® Corporation

PTP interaction with EPL

The following table describes the PTP packet disposition with respect to the configured PTP profile and the configured L2CFT control protocol (ptp-peer-delay) disposition over EPL (Ethernet Private Line).

Table 5-5
PTP interaction with EPL

Destination MAC Type	PTP(G.8275.1) Configured	L2CFT-PTP Disposition	Result
Non-Forwardable multicast MAC (01-80-C2-00-00-0E)	Yes	Discard	Peer
		Forward	Peer
	No	Discard	Discard
		Forward	Forward
Forwardable multicast MAC (01-1B-19-00-00-00)	Yes	Discard	Peer
		Forward	Peer
	No	Discard	Forward
		Forward	Forward

External timing

External timing is done by means of

- "BITS"
- "GPS"

BITS resides in the SYNC port.

BITS

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 in the transmitted data signal to synchronize the whole network.

The 3904, 3924, 3926, 3928, 5142, the optional 10G module with BITS for BITS support on 5160 platforms input and output standard electrical specifications for a BITS frequency as follows:

- 2.048 Mbps E1 compliant to ITU-T G.703-9
- 1.544 Mbps T1 compliant to GR-499 and ITU-T G.703-5
- 2048 KHz

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

Copyright© 2022 Ciena® Corporation

GPS

The 3904, 3924, 3926, 3928, 5142, and 5160 platforms can synchronize to a GPS receiver by means of:

- 10 MHz mini-coax connector, which allows a frequency reference signal to be used as an input or generated as an output. The frequency of the signal can be configured as 10 MHz, 2.048 MHz, or 1.544 MHz.
- 1 PPS mini-coax connector, which allows a 1 PPS signal to be used as an input or generated as an output for phase synchronization. This interface is also capable of recovering and generating embedded ToD messages. The 1 PPS interface is unidirectional: it can be configured either as input or output but not both simultaneously.
- SYNC port, which allows a 1 PPS signal to be used as an input or generated as an output for phase synchronization. This interface is also capable of recovering and generating ToD messages. The SYNC port is unidirectional: it can be configured either as input or output but not both simultaneously.

For more information, refer to “External timing,” in *39XX/51XX Switches and Platforms Product Fundamentals*.

Clock selection algorithm

The clock selection algorithm selects the best available synchronization source from the nominated clock references added to a protection-group. When configured as an Option 1 (E1) clock it operates in revertive mode by default. When configured as an Option 2 (T1) clock it operates in non-revertive mode by default.

In revertive mode, the best reference is selected after a wait-to-restore period has elapsed when a better reference is introduced. The wait-to-restore timer only applies to revertive-mode.

In non-revertive mode, once a reference has been selected, the clock selection algorithm does not select a new reference even if the new reference is better. A new reference is selected only when the currently-selected reference fails.

The clock selection algorithm uses these criteria to select the clock source from a set of configured clock references:

- operator commands, that is, force switch and clear
- loss-of-service, loss-of-frame, alarm indication signal, hardware not set up, or clock out-of-frame
- optional provisioned clock reference override-priority

Copyright© 2022 Ciena® Corporation

- optional provisioned clock reference forced quality level, which overrides received quality level
- received quality level (SSM/clockclass)
- optional provisioned clock reference priority

The events that can trigger the algorithm are:

- CLI commands:
  - add references to or remove references from a protection-group
  - Input reference override-priority, priority, forced-ql, or ql-receive configuration changes
  - forced reference configuration changes
  - threshold-ql configuration changes
  - reversion-mode configuration changes
- Physical characteristics:
  - received quality level value changes or timeout
  - loss-of-service, loss-of-frame, alarm indication signal, hardware not set up, or clock out-of-frame by means of QL-FAILED changes

When no priority is configured among the references, for a tie-breaker case, for example, no override-priorities/priorities are configured and all references have the same quality level, this order of protocol preference is used to determine the preferred reference:

- GPS
- BITS
- SyncE
- PTP
- TDM

In the case of a tie between references of the same protocol, the interface identifier breaks the tie, for example, for SyncE port 1 is preferred over port 10 (the lower port is preferred) and for PTP master-ip-address 1.1.1.1 is preferred over 2.2.2.2 (the lower address is preferred).

The clock selection algorithm applies to three separate protection groups:

- frequency
- phase
- time-of-day

Copyright© 2022 Ciena® Corporation

Quality level value definition

In SyncE, quality level is encoded in SSMs and is used to determine the validity of the signal. If the received quality level of a clock source is below the user-configured minimum quality level threshold, the clock source is deemed invalid and does not participate in the clock selection algorithm. You can configure a forced-ql attribute that can override the received quality level.

The BITS interface carries an SSM value defined by ITU-T G.781. Depending on the configuration of the BITS interface type, that is, T1 or E1, the value of SSM and the acceptable received quality levels for Sync-E and BITS are defined in Table 5-6 through Table 5-9. All other received SSM values are considered invalid on input.

For PTP, quality level is extracted from clockclass. For more information, refer to Table 5-6 to Table 5-9. To be selectable, references that do not carry a quality level must be configured with a forced-ql to meet clock selection algorithm criteria.

Different types of inputs can have different scales, for example, SSM/clockclass, to measure quality level. These quality levels are mapped to a common scale with a preference score, where a lower number determines higher preference, based on which an input is preferred over another. It is important to note that there are separate quality level mappings for input and output. Once all inputs are mapped to the common scale and the best input is selected using clock selection algorithm criteria, the preference score of the selected reference is then mapped back to SSM/clockclass on output depending on the protocol of the output references.

This table shows the input quality level mapping for Option I.

Table 5-6
Option I input quality level mapping

From BITS E1 SSM	From SyncE SSM	From PTP clockclass (G.8265.1)	QL name	From QL preference
		80		1
		82		2
0x2	0x2	84	QL-PRC	3
		86		4
		88		5
0x4	0x4	90	QL-SSU-A	6
		92		7

Copyright© 2022 Ciena® Corporation

Table 5-6
Option I input quality level mapping

From BITS E1 SSM	From SyncE SSM	From PTP clockclass (G.8265.1)	QL name	From QL preference
		94		8
0x8	0x8	96	QL-SSU-B	9
		98		10
		100		11
		102		12
0xB	0xB	104	QL-SEC/EEC1	13
		106		14
		108		15
0xF or QL-INVx	0xF or QL-INVx	110	QL-DNU	16

This table shows the output quality level mapping for Option I.

Table 5-7
Option I output quality level mapping

From QL preference	QL name	To BITS T1 SSM	To SyncE SSM	To PTP clockclass (G.8265.1)
1	QL-PRC	0x2	0x2	80
2				82
3				84
4	QL-SSU-A	0x4	0x4	86
5				88
6				90
7	QL-SSU-B	0x8	0x8	92
8				94
9				96
10				98
11				100

Copyright© 2022 Ciena® Corporation

Table 5-7
Option I output quality level mapping

From QL preference	QL name	To BITS T1 SSM	To SyncE SSM	To PTP clockclass (G.8265.1)
12	QL-SEC/EEC1	0xB	0xB	102
13				104
14				106
15	QL-DNU	0xF	0xF	108
16				110

This table shows the input quality level mapping for Option II.

Table 5-8
Option II input quality level mapping

From BITS E1 SSM	From SyncE SSM	From PTP clockclass (G.8265.1)	QL name	To QL preference
0x04FF	0x1	80	QL-PRS	1
0x08FF	0x0	82	QL-STU	2
		84		3
0x0CFF	0x7	86	QL-ST2	4
		88		5
0x78FF	0x4	90	QL-TNC	6
		92		7
		94		8
		96		9
		98		10
0x7CFF	0xD	100	QL-ST3E	11
0x10FF	0xA	102	QL-ST3/EEC2	12
		104		13

Copyright© 2022 Ciena® Corporation

Table 5-8
Option II input quality level mapping

From BITS E1 SSM	From SyncE SSM	From PTP clockclass (G.8265.1)	QL name	To QL preference
0x22FF	0xC	106	QL-SMC	14
0x40FF	0xE	108	QL-PROV	15
0x30FF or QL-INVx	0xF or QL-INVx	110	QL-DUS	16

This table shows the output quality level mapping for Option II.

Table 5-9
Option II output quality level mapping

From QL preference	QL name	To BITS T1 SSM	To SyncE SSM	To PTP clockclass (G.8265.1)
1	QL-PRS	0x04FF	0x1	80
2	QL-STU	0x08FF	0x0	82
3	QL-PRS	0x04FF	0x1	84
4	QL-ST2	0x0CFF	0x7	86
5	QL-TNC	0x78FF	0x4	88
6				90
7	QL-ST3E	0x7CFF	0xD	92
8				94
9				96
10				98
11				100
12	QL-ST3/EEC2	0x10FF	0xA	102
13	QL-SMC	0x22FF	0xC	104
14				106
15	QL-PROV	0x40FF	0xE	108
16	QL-DUS	0xF	0xF	110

Copyright© 2022 Ciena® Corporation

Frequency, phase and time-of-day configuration rules

In synchronization, frequency, phase and time-of-day depend on each other. Time-of-day reference depends on phase reference. Phase reference depends on frequency reference. To prevent unwanted configurations, these rules must be followed:

- References cannot be added to the time-of-day protection-group if the phase protection-group is empty.
- References cannot be added to the phase protection-group if the frequency protection-group is empty
- References can be added at any time to the frequency protection-group.
- References can be removed at any time from the time-of-day protection-group.
- The last reference can only be removed from the phase protection-group if the time-of-day protection-group is empty.
- The last reference can only be removed from the frequency protection-group if the phase protection-group is empty.

The time-of day protection-group depends on the phase protection-group. The phase protection-group depends on the frequency protection-group. The Configuration rules table shows the four valid combinations at any time. Anything else is blocked by the CLI.

Table 5-10
Configuration rules

	Frequency	Phase	Time-of-day	Description
1	X			Frequency references are added while time-of-day and phase protection-groups are empty.
2	X	X		Phase references are added while the time-of-day protection-group is empty and the frequency protection-group is not empty.
3	X	X	X	Time-of-day references are added while frequency and phase protection-groups are not empty.
4				All protection-groups are empty.

Note: The selected phase reference is traceable to the same primary reference clock as the selected frequency reference. The selected time-of-day reference is also traceable to the same primary reference clock as the selected phase reference. If the currently selected frequency reference goes down, but is still added to the frequency protection-group, the phase and time-of-day protection-groups disregard the selection and go into freerun/holdover. Similarly, when the currently selected phase

Copyright© 2022 Ciena® Corporation

reference goes down, but is still added to the phase protection-group, the time-of-day protection-group disregards the selection and goes into freerun/holdover.

Holdover interval

The holdover interval gives you the option of selecting how you want to control the protection-group operational quality level during holdover. This is configured using the `sync set holdover-interval <indefinite|24 hrs>` command. The holdover and free-run quality levels are now the oscillator quality level of the device. Use this command to have the protection-group operational quality level transition to DNU/DUS (do not use) after 24 hours of being in holdover. To be consistent with G.781, the default holdover interval is indefinite.

PTP clock type

The PTP clock type can be configured as ordinary clock slave or boundary clock. You can set the PTP clock type from ordinary clock slave to boundary clock or vice versa.

When the PTP clock-type is set from ordinary clock slave to boundary clock, the CLI prints a message informing you to remove all GPS references that were added to the protection-groups. When the PTP clock-type is set to boundary clock, you can then create PTP output references, and you cannot add any GPS references to protection-groups.

PTP outputs can only be derived from non-GPS inputs when PTP clock type is set to boundary clock. In boundary clock mode, GPS input references are disregarded for selection. Only non-GPS inputs qualify for selection while timing distribution can occur via both GPS and non-GPS outputs.

If you set PTP clock-type from boundary clock to ordinary clock slave, the CLI prints a message informing you to delete all PTP output references. When the PTP clock is set to ordinary clock slave, you are not able to create PTP output references, and you can add any GPS references to protection-groups.

Reversion-mode configuration is ignored while switching the PTP clock-type.

Copyright© 2022 Ciena® Corporation

Network configuration examples

The Linear topology figure 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 5-4
Linear topology

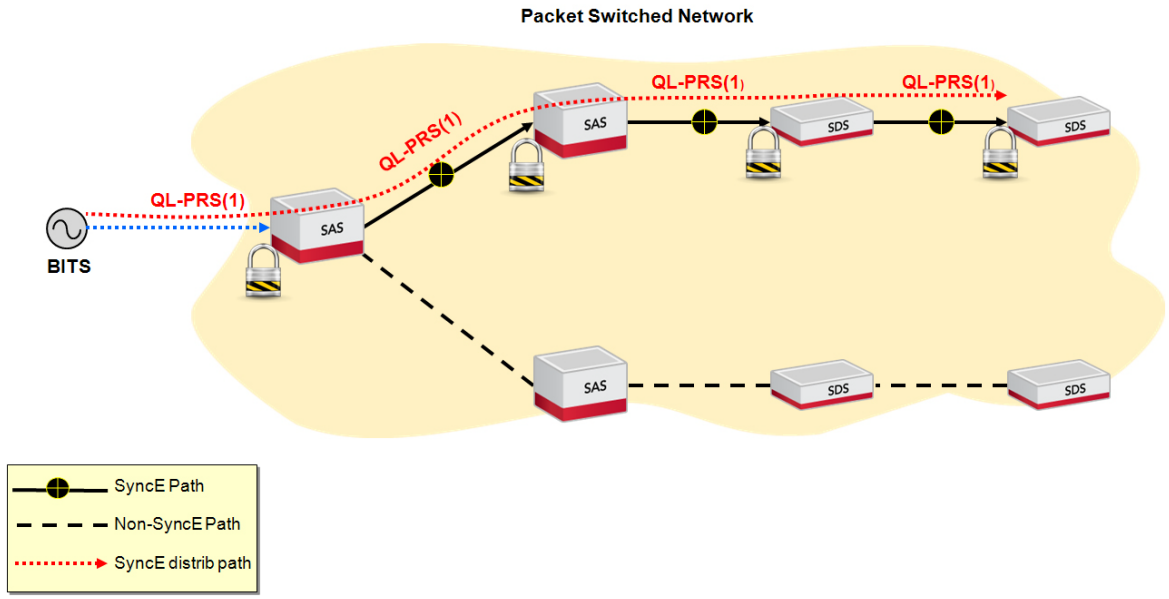

Copyright© 2022 Ciena® Corporation

The Ring topology figure 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 5-5
Ring topology

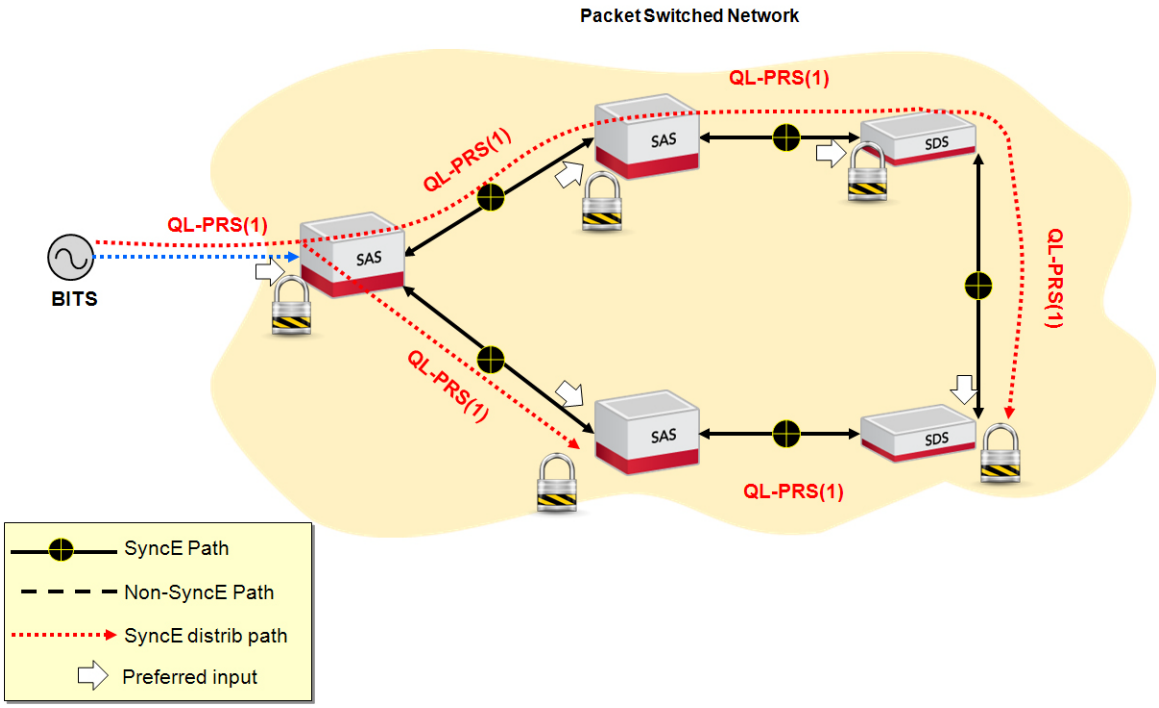

Copyright© 2022 Ciena® Corporation

The System timing using SyncE and PTP figure shows an example of system timing using SyncE and PTP. In this example, SyncE is used to recover frequency and PTP is used to recover phase.

Figure 5-6
System timing using SyncE and PTP

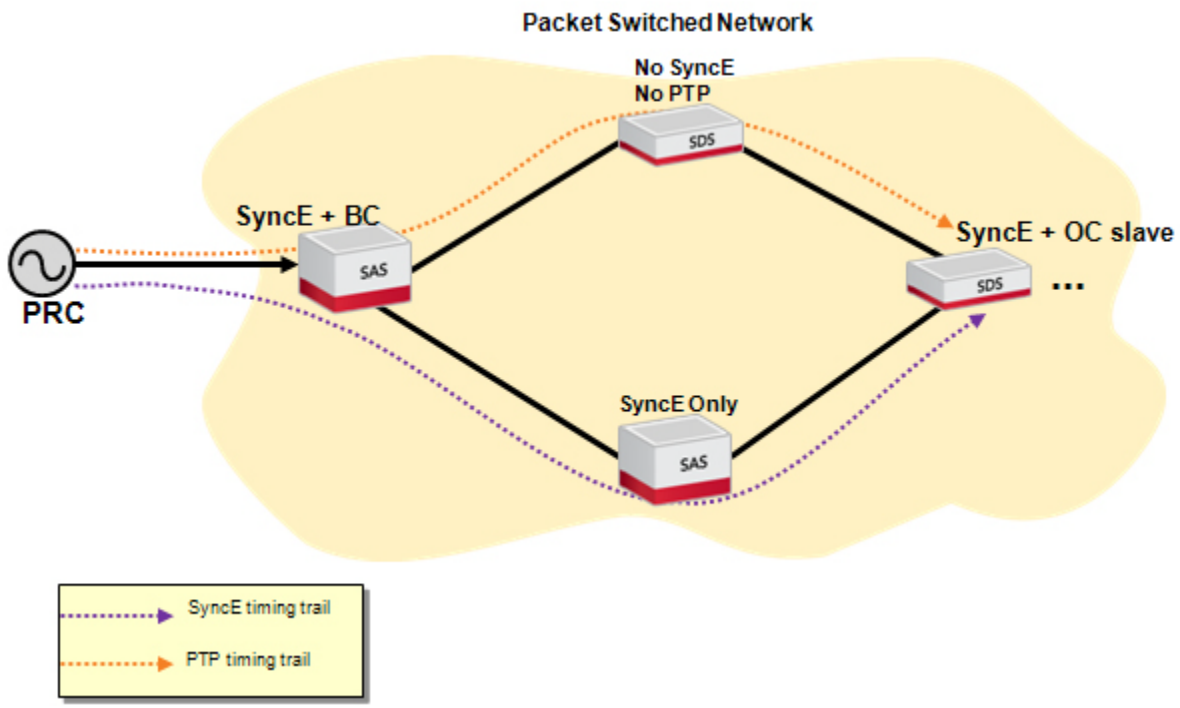

"Sample configuration: system timing by means of SyncE and PTP" on page 5-91 provides the commands to configure this example.

Copyright© 2022 Ciena® Corporation

This figure shows an example of recovering PTP timing on an Ordinary Clock Slave from a GrandMaster Clock via a Boundary Clock (priority 1) while keeping the GrandMaster Clock as a backup timing reference (priority 2).

Figure 5-7
System timing using PTP timing

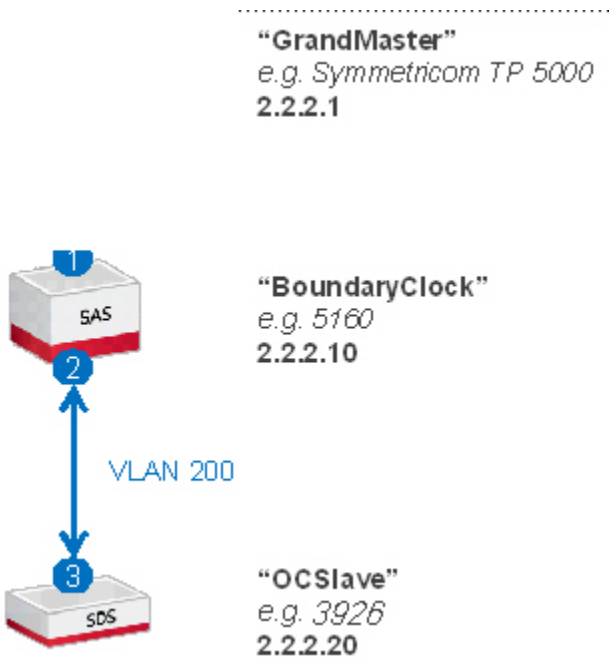

"Sample configuration: system timing by means of PTP boundary clock" on page 5-93 provides the commands to configure this example.

Copyright© 2022 Ciena® Corporation

Procedures

This figure shows the flow of system timing procedures.

Figure 5-8
System timing configuration

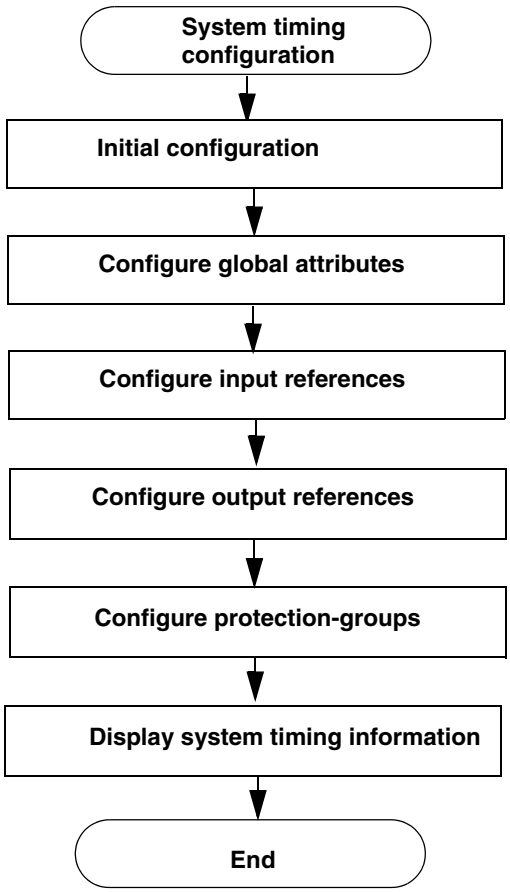

Note: If the AE and AOAM license key is not installed, system timing commands are disabled.

Procedures for initial configuration are:

- "Enabling and disabling synchronization" on page 5-25
- "Configuring synchronization" on page 5-26

Procedures for configuring the global attributes for GPS output timing

- "Configuring the global PTP telecom profile" on page 5-29
- "Configuring the PTP timing global attributes" on page 5-30
- "Configuring global attributes for PTP input timing" on page 5-32

Copyright© 2022 Ciena® Corporation

- “Configuring global attributes for PTP output timing” on page 5-33
- “Configuring global attributes for GPS output timing” on page 5-35

Procedure for configuring input references are:

- “Configuring SyncE input references” on page 5-36
- “Configuring BITS input references” on page 5-38
- “Configuring PTP input references” on page 5-41
- “Configuring GPS input references” on page 5-44

Procedures for configuring output references are:

- “Configuring SyncE output references” on page 5-46
- “Configuring BITS output reference” on page 5-47
- “Configuring PTP output timing references” on page 5-49
- “Configuring GPS output references” on page 5-51

Procedure for configuring protection-groups are:

- “Configuring protection-groups” on page 5-52

Procedures for displaying system timing information are:

- “Displaying information for synchronization” on page 5-61
- “Displaying SyncE information” on page 5-65
- “Displaying BITS information” on page 5-68
- “Displaying PTP information” on page 5-71
- “Displaying GPS information” on page 5-78
- “Displaying TDM information” on page 5-81
- “Displaying frequency information” on page 5-83
- “Displaying phase information” on page 5-85
- “Displaying time-of-day information” on page 5-87
- “Displaying protection-group information” on page 5-89

Additionally, you can clear timing statistics. See “Clearing timing statistics” on page 5-60.

For information about commands used to unset system timing configurations, refer to *39XX/51XX Switches and Platforms Command Reference*.

Copyright© 2022 Ciena® Corporation

Procedure 5-1
Enabling and disabling synchronization

You can enable and disable synchronization. Synchronization is enabled by default.

Disabling synchronization results in:

- the clock selection algorithm is frozen, that is, the system no longer responds to events that trigger the clock selection algorithm
- output references are shut down so that downstream nodes no longer receive synchronization. Prior to shutting down the output references, the output references that are capable of transmitting quality level transmit QL-DNU/QL-DUS (do-not-use) messages to notify downstream nodes to gracefully switch over to another reference.

Disable the synchronization feature for maintenance purposes.

Step	Action
To enable synchronization	
1	Enable synchronization: `sync enable`
To disable synchronization	
2	Disable synchronization: `sync disable`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 5-2
Configuring synchronization

You can configure

- hold-over interval period, which can be configured to be 24-hrs or indefinite.
- option type, which is a one-time mandatory configuration before all other sync inputs, outputs, or protection-groups can be configured
- reversion mode, which affects the behavior of clock selection algorithm
- global wait-to-restore timer, which is related to reversion mode. The wait to restore timer applies when the reversion mode is set to revertive.

Step	Action

To set the holdover interval

1 Set the holdover interval:

```
sync set holdover-interval <indefinite | 24-hrs>
```

where

holdover-interval <indefinite | 24-hrs>	• is the holdover interval. The default value is indefinite. • If configured to be 24-hrs, the operational QL of the protection-group changes from local oscillator QL to QL-DUS/QL-DNU after being in holdover for 24 hours. • If the hold-over interval is configured to be indefinite, the operational QL of the protection-group does not change during holdover.

Copyright© 2022 Ciena® Corporation

To set the option type

2 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 1:*** When option-type is set to option1, reversion-mode defaults to revertive. When option-type is set to option2, reversion-mode defaults to non-revertive. You can explicitly configure reversion-mode. ***Note 2:*** Selection of the type determines the quality level definition. For option1 quality level definitions, refer to the Option I input quality level mapping and Option I output quality level mapping tables. For option2 quality level definitions, refer to the Option II input quality level mapping and Option II output quality level mapping tables. To change the option-type, make sure there are no existing configurations.

To set the reversion-mode

3 Set the global reversion-mode:

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

where

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

To set the global wait-to-restore timer

4 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.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example sets the holdover-interval to 24-hrs.

```
sync set holdover-interval 24-hrs
```

This example sets the option type to option 1, for an SDH network optimized for the 2048 kbit/s hierarchy. Note that the reversion mode defaults to revertive.

```
sync set option-type option1
```

This example sets the global reversion mode to non-revertive.

```
sync set reversion-mode non-revertive
```


Copyright© 2022 Ciena® Corporation

Procedure 5-3
Configuring the global PTP telecom profile

You can configure the global PTP telecom profile for:

- G.8265.1—frequency synchronization
- G.8275.1—phase and time synchronization

You must delete all existing PTP references and disable synchronization before configuring the global PTP telecom profile.

Step	Action

To configure the global PTP telecom profile

1 Configure the global PTP telecom profile:

```
sync ptp set profile <none | g.8265.1 | g.8275.1>
```

where

profile <none | g.8265.1 | g.8275.1>	is the global PTP telecom profile. The default value is none. • g.8265.1 is the PTP telecom profile for frequency synchronization. • g.8275.1 is the PTP telecom profile for phase and time synchronization.

—end—

Example

This example sets the global PTP telecom profile to G.8275.1.

```
sync ptp set profile g.8275.1
```


Copyright© 2022 Ciena® Corporation

Procedure 5-4
Configuring the PTP timing global attributes

You can configure these PTP timing attributes:

- address mode
- domain number
- profile identifier
- profile version
- tag priority
- clock type
- dscp value
- destination mac

Step	Action
1	Globally set PTP timing attributes for the network element: `sync ptp set {address-mode <unicast>} {domain-number <NUMBER: 4-23>} {profile-identifier <00-19-A7-00-01-00>} {profile-version <1.0>} {protocol-version <2>} {tag-priority <NUMBER: 0-7>} clock-type <oc-slave|bc> dscp <0...63> destination-mac <forwardable|non-forwardable>` where address-mode <unicast> — is the PTP address mode. The default value is `unicast`. domain-number <NUMBER: 4-23> — is the PTP domain number. profile-identifier <00-19-A7-00-01-00> — is the profile identifier for the network element. profile-version <1.0> — is the profile version supported by the network element. protocol-version <2> — is the PTP version supported by the network element. tag-priority <NUMBER: 0..7> — is the VLAN p-bits for outgoing PTP packets. The default value is `0`.

Copyright© 2022 Ciena® Corporation

where

clock-type	is the PTP clock type. The default is `oc-slave`.
dscp <0..63>	sets the dscp value for PTP traffic.
destination-mac	configures the multicast MAC address to allow sessions to transmit forwardable/non-forwardable addresses on all PTP ports. The MAC address is selected automatically when one of the following options is selected: • forwardable - `01:1b:19:00:00:00` • non-forwardable - `01:80:c2:00:00:0e`

—end—

Example

This example sets global PTP timing attributes for the network element.

```
sync ptp set address-mode unicast domain number 4 profile-
identifier 00-10-A7-00-01-00 profile-version 1.0 protocol-
version 2 tag-priority 0 clock-type bc dscp 56
```

This example configures the multicast MAC address as forwardable.

```
sync ptp set destination-mac forwardable
```

This example configures the multicast MAC address as non-forwardable.

```
sync ptp set destination-mac non-forwardable
```


Copyright© 2022 Ciena® Corporation

Procedure 5-5
Configuring global attributes for PTP input timing

Configure global attributes for PTP input timing.

Step	Action
1	Set global attributes for PTP input timing: `sync ptp input set {announce-rx-loss-num <NUMBER: 2..10>} {announce-tx-rate-request <1|1-2nd|1-4th|1-8th>} {delay-request-tx-rate <64|32|16|8|4|2|1|1-2nd>} {sync-tx-rate-request <64|32|16|8|4|2|1|1-2nd>} {unicast-request-duration <SECONDS: 60-1000>}` where announce-rx-loss-num <NUMBER: 2..10> — is the PTP input announce receive loss number (announce intervals). The default value is 3. announce-tx-rate-request <1|1-2nd|1-4th|1-8th> — is the PTP input announce transmit rate request, in packets per second. The default value is 1-2nd. delay-request-tx-rate <64|32|16|8|4|2|1|1-2nd> — is the PTP input delay request transmit rate, in packets per second. The default value is 64. sync-tx-rate-request <64|32|16|8|4|2|1|1-2nd> — is the PTP input sync transmit rate request, in packets per second. The default value is 64. unicast-request-duration <SECONDS: 60-1000> — is the PTP input unicast request duration, in seconds. The default value is 300. **—end—**

Example

This example sets global attributes.

```
sync ptp input set unicast-request-duration 300 announce-rx-
loss-num 3 announce-tx-rate-request 1-2nd sync-tx-rate-request
64 delay-request-tx-rate 64
```


Copyright© 2022 Ciena® Corporation

Procedure 5-6
Configuring global attributes for PTP output timing

The maximum number of clients (slave sessions) is configurable from 0 to 16. The default value is 16. Each platform has its limitation as shown in this table.

Table 5-11
Recommended values for max-slave-sessions by platform

Platform	Recommended max-slave-sessions
3904	6
3905	6
3924	6
3926	6
3928	6
5142	16
5160	16
Note: These numbers are recommendations only. You can connect as many slave sessions as you want, but going beyond these recommendations may lead to unacceptable clock recovery performance.	

Configure global attributes for PTP output timing based on the platform's recommended max-slave-session values.

Copyright© 2022 Ciena® Corporation

Step	Action
1	Set global attributes for PTP output timing: `sync ptp output set {max-slave-sessions <NUMBER 0..16>} {timestamp-mode <one-step|two-step>}` where max-slave-sessions <NUMBER 0-16> — is the maximum number of slave sessions timestamp-mode <one-step|two-step> — sets the timestamp mode to be either one-step or two-step ***Note:*** One-step timestamping is supported on the 5142 and 5160 only. **—end—**

Example

This example sets global attributes for PTP output timing.

```
sync ptp output set max-slave sessions 15 timestamp-mode one-
step
```


Copyright© 2022 Ciena® Corporation

Procedure 5-7
Configuring global attributes for GPS output timing

Configure global attributes for GPS output timing.

Step	Action
1	Set the 1 PPS pulse width: `sync gps output set {1pps-pulse-width <MICROSECONDS>}` where 1pps-pulse-width <MICROSECONDS> is the GPS output 1 PPS pulse width in microseconds.

—end—

Example

This example sets 1 PPS pulse width.

```
sync gps output set 1 pps-pulse-width 1000
```


Copyright© 2022 Ciena® Corporation

Procedure 5-8
Configuring SyncE input references

Configure SyncE input references to recover frequency timing from Ethernet ports.

You can:

- create a SyncE input reference
- set attributes for a SyncE input reference

Step	Action

To create a SyncE input reference

1 Create a SyncE input reference:

```
sync synce input create ref <ref> {forced-ql <Clock
Quality Level>} {override-priority <NUMBER: 1-10>} {port
<Physical port list>} {priority <NUMBER: 1-10>} {ql-
receive <disable|enable>}
```

where

ref <ref>	is the identifier for the SyncE 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, with 1 being the highest override priority.
port <phy-port-name>	is the physical port.
priority <NUMBER: 1-10>	is the priority, with 1 being the highest priority.
ql-receive disable|enable	enables or disables received quality level.

Copyright© 2022 Ciena® Corporation

To set attributes for a SyncE input reference

2 Set attributes for a SyncE input reference:

```
sync synce input set ref <ref> {forced-ql <Clock Quality
Level>} {priority <NUMBER: 1-10>} {override-priority
<NUMBER: 1-10>} {ql-receive <disable|enable>}
```

where

ref <ref>	is the identifier for the SyncE input reference.
forced-ql <INTEGER: 0-15>	sets the quality level at which the received quality level is overridden.
priority <NUMBER: 1-10>	is the priority, with 1 being the highest priority.
override-priority <INTEGER: 1-10>	is the override-priority, with 1 being the highest override priority.
ql-receive <disable|enable>	enables or disables received quality level.

—end—

Example

This example creates a SyncE input reference named mySyncEinput5 on port 5 with a priority of 1.

```
sync synce input create ref SyncE_Port5_In port 5 priority 1
```


Copyright© 2022 Ciena® Corporation

Procedure 5-9
Configuring BITS input references

Configure BITS input references to recover frequency timing from the SYNC interface.

You can:

- create a BITS input reference
- set attributes for a BITS input reference

Step	Action

To create a BITS input reference

1 Create a BITS input reference:

```
sync bits input create ref <ref> {bits-interface <BITS
interface>} {bits-mode <e1|t1|2048k>} {encoding
<ami|b8zs|hdb3>} {e1-ssm-location <sa4|sa5|sa6|sa7>}
{forced-ql <Quality Level>} {format <esf|e1-crc|e1-no-
crc|sf>} {override-priority <NUMBER: 1-10>} {priority
<NUMBER: 1-10>} {ql-receive <disable|enable>}
```

where

ref <ref>	is the BITS input reference.
bits-interface <BITS interface>	is the BITS interface.
bits-mode <e1|t1|2048k>	is the mode.
encoding <ami|b8zs|hdb3>	is the encoding type. For T1 mode, the default value is ami. For E1 mode, the default value is hdb3.
e1-ssm-location <sa4|sa5|sa6|sa7>	is the location for SSM messaging. This attribute applies only for E1 mode. The default value is sa4.
forced-ql <Quality Level>	is the quality level at which the received quality level is overridden.
format <esf |e1-crc| e1-no-crc|sf>	is the format.

Copyright© 2022 Ciena® Corporation

where

override-priority <NUMBER: 1..10>	is the override-priority, with 1 being the highest override priority.
priority <NUMBER: 1..10>	is the priority, with 1 being the highest priority.
ql-receive <disable|enable>	enables or disables received quality level. For E1 mode and format e1-crc, the default value is enable. For T1 mode and format esf, the default value is enable.

To set attributes for a BITS input reference

2 Set attributes for a BITS input reference:

```
sync bits input set ref <ref> {encoding <ami|b8zs|hdb3>}
{e1-ssm-location <sa4|sa5|sa6|sa7>} {forced-ql <Quality
Level>} {override-priority <NUMBER: 1-10>} {priority
<NUMBER: 1-10>} {ql-receive <disable|enable>}
```

where

ref <ref>	is the BITS input reference.
encoding <ami|b8zs|hdb3>	is the encoding type. For T1 mode, the default value is ami. For E1 mode, the default value is hdb3.
e1-ssm-location <sa4|sa5|sa6|sa7>	is the location for SSM messaging. This attribute applies only for E1 mode. The default value is sa4.
forced-ql <Quality Level>	is the quality level at which the received quality level is overridden.
override-priority <NUMBER: 1..10>	is the override-priority, with 1 being the highest override priority.
priority <NUMBER: 1..10>	is the priority, with 1 being the highest priority.
ql-receive <disable|enable>	enables or disables received quality level. For E1 mode and format e1-crc, the default value is enable. For T1 mode and format t1-esf, the default value is enable.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example creates a BITS input reference named myBITSinput1 with a BITS interface of sync-rj-45-1 and a mode of T1.

```
sync bits input create ref myBITSinput1 bits-interface sync-
rj45-1 bits-mode t1 format esf
```


Copyright© 2022 Ciena® Corporation

Procedure 5-10
Configuring PTP input references

Configure PTP input references to synchronize timing, that is, frequency (G.8265.1), phase and time-of-day (G.8275.1), from a grandmaster or boundary clock over a packet network.

You can

- create PTP input timing references
- set attributes for a PTP input timing reference

Note: PTP behavior over link aggregation groups (LAG)—Currently PTP is supported over LAG limited to a single port member. In the case where LAG is used for a timing reference, you cannot add a second member port to the LAG. In the case where a LAG has two member ports, you cannot create PTP over this LAG.

Step	Action

To create a PTP input timing reference for G.8265.1

1 Create a PTP input timing reference:

```
sync ptp input create ref <ref> {forced-ql <Clock Quality
level>} {ip-interface <Interface Name>} {master-ip-
address <IP address>} {override-priority <NUMBER: 1-255>}
{priority <NUMBER: 1-255>} {forced-clock-class <NUMBER:
80-110>} {ql-receive <disable | enable>}
```

where

ref <ref>	is the PTP input timing reference.
forced-ql <Clock Quality level>	is the forced quality level.
ip-interface <Interface Name>	is the IP interface name.
master-ip-address <IP address>	is the master IP address.
override-priority <NUMBER: 1-255>	is the override priority, with 1 being the highest override priority.

Copyright© 2022 Ciena® Corporation

where

priority <NUMBER: 1-255>	is the priority, with 1 being the highest priority.
forced-clock-class	is the forced clock class from 80-110.
ql-receive <disable | enable>	is the quality level receiving configuration.

To create a PTP input timing reference for G.8275.1

2 Create a PTP input timing reference:

```
sync ptp input create ref <ref> port <port> {override-
priority <NUMBER: 1-255>}{priority <NUMBER: 1-255>}
```

where

ref <ref>	is the PTP input timing reference.
port <port>	is the physical port.
override-priority <NUMBER: 1..255>	is the override priority, with 1 being the highest override priority.
priority <NUMBER: 1..255>	is the priority, with 1 being the highest priority.

To set attributes for a PTP input timing reference for G.8265.1

3 Set attributes for a PTP input timing reference:

```
sync ptp input set ref <ref> {forced-ql <Clock Quality
level>} {override-priority <NUMBER: 1-255>} {priority
<NUMBER: 1-255>} {forced-clock-class <80-110>} {ql-
receive <disable|enable>}
```

where

ref <ref>	is the PTP input timing reference.
forced-ql <Clock Quality level>	is the forced quality level.
override-priority <NUMBER: 1-255>	is the override priority, with 1 being the highest override priority.

Copyright© 2022 Ciena® Corporation

where

priority <NUMBER: 1-255>	is the priority, with 1 being the highest priority.
forced-clock-class	Is the forced clock class from 80-110.
ql-receive <disable | enable>	is the quality level receiving configuration.

To set attributes for a PTP input timing reference for G.8275.1

4 To set attributes for a PTP input timing reference:

```
sync ptp input set ref <ref> {override-priority <NUMBER:
1-255>} {priority <NUMBER: 1-255>}
```

where

ref <ref>	is the PTP input timing reference.
override-priority <NUMBER: 1-255>	is the override priority, with 1 being the highest override priority.
priority <NUMBER: 1-255>	is the priority, with 1 being the highest priority.

—end—

Examples

This example creates a PTP input timing reference for frequency named PTP_In.

```
sync ptp input create ref PTP_In master-ip-address 2.2.2.1 ip-
interface PTP_Interface priority 2
```

This example sets attributes for a PTP input timing reference for frequency named PTP_In.

```
sync ptp input set ref PTP_In priority 2 forced-clock-class 80
ql-receive enable
```

This example creates a PTP input timing reference for phase and time of day named PTP_In_2 on port1.

```
sync ptp input create ref ptp_In_2 port port1
```

This example sets attributes for a PTP input timing reference for phase and time of day named PTP_In_1.

```
sync ptp input set ref ptp_In_1 priority 2
```


Copyright© 2022 Ciena® Corporation

Procedure 5-11
Configuring GPS input references

Configure GPS input references to recover timing from a GPS signal. GPS input references can be configured by means of

- 10MHz coax connection (frequency only)
- 1PPS coax connection (phase and time-of-day)
- SYNC interface (phase and time-of-day)

You can

- create a GPS input reference
- set attributes of a GPS input reference

Step	Action

To create a GPS input reference

1 Create a GPS input reference:

```
sync gps input create ref <ref> {clock-mode
<frequency|phase>} {forced-ql <Clock Quality Level>}
{frequency-clock <10mhz|1544khz|2048khz>} {gps-interface
<GPS interface>} {override-priority <NUMBER: 1-10>}
{priority <NUMBER: 1-10>} frequency-clock <10mhz|2048
khz|1544khz>
```

where

ref <ref>	is the GPS timing interface.
clock-mode <frequency| phase>	is the GPS clock mode.
forced-ql <Clock Quality level>	is the forced quality level.
frequency-clock <10mhz| 1544khz| 2048khz>	is the GPS frequency clock.

Copyright© 2022 Ciena® Corporation

where

gps-interface <GPS interface>	is the GPS interface.
override-priority <NUMBER: 1-10>	is the override priority, with 1 being the highest override priority.
priority <NUMBER: 1-10>	is the priority, with 1 being the highest priority.

To set attributes of a GPS input reference

2 Set attributes of a GPS input reference:

```
sync gps input set ref <ref-name> {forced-ql <Clock
Quality Level>} {override-priority <NUMBER: 1-10>}
{priority <NUMBER: 1-10>}
```

where

ref <ref>	is the GPS input timing reference.
forced-ql <Clock Quality level>	is the forced quality level.
{override-priority <NUMBER: 1-10>	is the override priority, with 1 being the highest override priority.
priority <NUMBER: 1-10>	is the priority, with 1 being the highest priority.

—end—

Example

This example creates a GPS input reference named myGPS10MHzInput1.

```
sync gps input create ref myGPS10MHzInput1 clock-mode frequency
forced-ql prs frequency-clock 10mhz gps-interface 10mhz-smb-1
```


Copyright© 2022 Ciena® Corporation

Procedure 5-12
Configuring SyncE output references

Configure SyncE output references to distribute frequency timing by means of Ethernet ports.

> ***Note:*** On the 3924, 3926 and 3928 platforms, ports 3 and 4 do not support SyncE output at the 10Gig rate.

Step	Action
1	Create a SyncE output reference: `sync synce output create ref <ref> port <Physical port>` where ref <ref> — is the SyncE output timing reference. port <Physical port list> — is the physical port. **—end—**

Example

This example creates a SyncE output reference named mySyncEoutput1 on port 1.

```
sync synce output create ref mySyncEoutput1 port 1
```


Copyright© 2022 Ciena® Corporation

Procedure 5-13
Configuring BITS output reference

Configure BITS output references to distribute frequency timing by means of the SYNC interface.

You can

- create a BITS output reference
- set attributes for a BITS output reference

Step	Action

To create a BITS output reference

1 Create a BITS output reference:

```
sync bits output create ref <ref> {bits-interface <BITS
interface>} {bits-mode <e1|t1|2048k>} {encoding
<ami|b8zs|hdb3>} {e1-ssm-location <sa4|sa5|sa6|sa7>}
{format <e1-crc|e1-no-crc|esf|sf>} {t1-line-build-out
<133|266|399|533|655>}
```

where

ref <ref>	is the BITS output reference.
bits-interface <BITS interface>	is the BITS interface.
bits-mode <e1|t1|2048k>	is the mode.
encoding <ami|b8zs|hdb3>	is the encoding type. For T1 mode, the default value is ami. For E1 mode, the default value is hdb3.
e1-ssm-location <sa4|sa5|sa6|sa7>	is the location for SSM messaging. This attribute applies only for E1 mode. The default value is sa4.
format <e1-crc|e1-no-crc|t1-esf|t1-sf>	is the format.
t1-line-build-out <133|266|399|533|655>	is the BITS 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. The default value is 133.

Copyright© 2022 Ciena® Corporation

To set attributes for a BITS output reference

2 Set attributes for a BITS output reference:

```
sync bits output set ref <ref> {encoding <ami|b8zs|hdb3>}
{t1-line-build-out <133|266|399|533|655>} {e1-ssm-
location <sa4|sa5|sa6|sa7>}
```

where

ref <ref>	is the BITS input reference.
encoding <ami|b8zs|hdb3>	is the encoding type. For T1 mode, the default value is ami. For E1 mode, the default value is hdb3.
t1-line-build-out <133|266|399|533|655>	is the BITS 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. The default value is 133.
e1-ssm-location <sa4|sa5|sa6|sa7>	is the location for SSM messaging. This attribute applies only for E1 mode. The default value is sa4.

—end—

Example

This example creates a BITS output reference named E1_BITS_Out.

```
sync bits output create ref E1_BITS_Out bits-interface sync-
rj45-1 bits-mode e1 format e1-crc
```


Copyright© 2022 Ciena® Corporation

Procedure 5-14
Configuring PTP output timing references

Configure PTP output references to distribute frequency (G.8265.1), phase and time-of-day (G.8275.1) information to connected PTP clients. You need to set the PTP clock-type to “bc” (boundary clock) to configure PTP output references (see “Configuring the PTP timing global attributes” on page 5-30).

Note: PTP behavior over link aggregation groups (LAG)—Currently PTP is supported over LAG limited to a single port member. In the case where LAG is used for a timing reference, you cannot add a second member port to the LAG. In the case where a LAG has two member ports, you cannot create PTP over this LAG.

Step	Action

To create a PTP output timing reference for G.8265.1

1 Create a PTP output reference:

```
sync ptp output create ref <ref-name> ip-interface <IP
interface>
```

where

<ref-name>	is the PTP output reference being created.
ip-interface <IP interface>	is the IP interface.

To create a PTP output timing reference for G.8275.1

2 Create a PTP output reference:

```
sync ptp output create ref <ref-name> port <port>
```

where

<ref-name>	is the PTP output reference being created.
port <port>	is the physical port number.

Copyright© 2022 Ciena® Corporation

To set a PTP output timing reference for G.8275.1

3 Set a PTP output reference:

```
sync ptp output set ref <ref-name> priority2 <NUMBER:
0...255> timestamp-mode <one-step | two-step> announce-
tx-rate <1 | 1-2nd | 1-4th | 1-8th | 2 |4 |8> sync-tx-rate
<1 | 1-2nd | 1-4th | 1-8th | 1-16th | 128 | 2 |4 |8 |16 |
32 | 64>
```

where

<ref-name>	is the PTP output reference being modified.
priority2 <NUMBER: 0...255>	is the priority of the PTP output reference.
timestamp-mode <one-step | two-step	is the timestamp mode.
announce-tx-rate `<1 | 1-2nd | 1-4th | 1-8th | 2 |4 |8>`	is the announcement transmission rate.
sync-tx-rate `<1 | 1-2nd | 1-4th | 1-8th | 1-16th | 128 | 2 |4 |8 |16 | 32 | 64>`	is the synchronization transmission rate.

—end—

Examples

This example creates a PTP output reference for frequency.

```
sync ptp output create ref PTPOut2 ip-interface PTP_Interface
encap-type udp-over-ip
```

This example creates a PTP output reference for phase and time of day.

```
sync ptp output create ref PTPOut1 port 1
```

This example sets a PTP output reference for phase and time of day.

```
sync ptp output set ref PTPOut3 priority2 10
```


Copyright© 2022 Ciena® Corporation

Procedure 5-15
Configuring GPS output references

Configure GPS output references to distribute timing to a neighboring device that accepts a GPS signal. GPS output references can be configured by means of

- 10MHz coax connection (frequency only)
- 1PPS coax connection (phase and time-of-day)
- SYNC interface (phase and time-of-day)

Step	Action
1	Create a GPS output reference: `sync gps output create ref <ref> {clock-mode <frequency|phase>} {frequency-clock <10mhz|2048khz|1544khz>} {gps-interface <GPS interface>}` where ref <ref> — is the GPS output timing reference. clock-mode <frequency | phase> — is the GPS clock mode. frequency-clock <10mhz| 1544khz| 2048khz> — is the GPS frequency clock. gps-interface <GPS interface> — is the GPS interface. **—end—**

Example

This example creates a GPS output reference named SMB_2048_Out.

```
sync gps output create ref SMB_2048_Out gps-interface 10mhz-
smb-1 clock-mode frequency frequency-clock 2048khz
```


Copyright© 2022 Ciena® Corporation

Procedure 5-16
Configuring protection-groups

Create protection-groups to initiate a clock selection algorithm session for the type of timing, that is, frequency or phase. Only input references that are added to a protection-group can take part in the clock selection process.

For Release 6.13, time-of-day extracted from PTP is only displayed. Full time-of-day functionality will be supported in a future release. In Release 6.13, PTP outputs can only be derived from non-GPS inputs when PTP clock-type is set to “boundary clock”. In boundary clock mode, GPS input references are disregarded for selection and only non-GPS inputs are qualified for selection while timing distribution can take place via GPS and non-GPS outputs.

This table list the maximum number of input references for each protection-group by platform.

Table 5-12
Maximum input references for each protection-group

Platform	Frequency					Phase		Time-of-Day+	
	SyncE	PTP	BITS	GPS	TDM	PTP	GPS	PTP+	GPS+
3904	4	6	1	1	0	6	2	6+	2+
3905	4	6	0	0	0	6	0	6+	0+
3924	8	6	1	1	2		1	6+	1
3926	8	6	1	1	2		1	6+	1
3928	12	6	1	1	0		1	6+	1
5142	24	6	1	1	0	6	2	6+	2+
5160	24	6	1	1	0	6	2	6+	2+

Note 1: + Although time-of-day protection-group is configurable, recovering time-of-day is not yet supported. Currently only PTP input references can be added to a time-of-day protection-group. The time-of-day protection-group selects the best PTP input reference and only displays its time-of-day information in the `sync time-of-day protection-group show <group-object>` command.

Note 2: It is important to make sure the selected time-of-day reference is traceable to the same selected phase reference and the selected phase reference is traceable to the same selected frequency reference.

Copyright© 2022 Ciena® Corporation

The color of the SYNC interface LED depends on the state of all configured protection-groups. This table lists protection-group states and describes the SYNC interface LED associated with each protection-group state.

Table 5-13
Protection-group states

Protection-group states	SYNC interface LED
There are no protection-groups configured for frequency, phase, and time-of-day.	Off
One or more configured frequency, phase, or time-of-day protection-groups is in Free Run state.	Off
One or more configured frequency, phase, or time-of-day protection-groups is in Holdover state, and all other configured protection-groups are in Holdover, Acquiring Lock, or Locked state.	Solid yellow
One or more configured frequency, phase, and time-of-day protection-groups is in Acquiring Lock state, and all other configured protection-groups are in Acquiring Lock or Locked state.	Blinking green
All configured frequency, phase, and time-of-day protection-groups are in Locked state.	Solid green

You can

- create a frequency protection-group
- create a phase protection-group
- create a time-of-day protection-group
- set attributes for a frequency protection-group
- set attributes for a phase protection-group
- set attributes for a time-of-day protection-group
- add an input reference to a frequency protection-group
- add an input reference to a phase protection-group
- add an input reference to a time-of-day protection-group
- remove an input reference from a frequency protection-group
- remove an input reference from a phase protection-group
- remove an input reference from a time-of-day protection-group
- clear a frequency protection-group's reference switch count
- clear a phase protection-group's reference switch count

Copyright© 2022 Ciena® Corporation

- clear a time-of-day protection-group's reference switch count

Step	Action

To create frequency protection-groups

1 Create a frequency protection-group:

```
sync frequency protection-group create group <group>
{threshold-ql <clock quality level>}
```

where

group <group>	is the frequency protection-group.
threshold-ql <clock quality level>	is the threshold quality level.

To create phase protection-groups

2 Create a phase protection-group:

```
sync phase protection-group create group <group>
{threshold-ql <clock quality level>}
```

where

group <group>	is the phase protection-group.
threshold-ql <clock quality level>	is the threshold quality level.

To create time-of-day protection-groups

3 Create a time-of-day protection-group:

```
sync time-of-day protection-group create group <group>
{threshold-ql <clock quality level>}
```

where

group <group>	is the time-of-day protection-group.
threshold-ql <clock quality level>	is the threshold quality level.

Copyright© 2022 Ciena® Corporation

To set attributes for a frequency protection-group'

4 Set attributes for a frequency protection-group:

```
sync frequency protection-group set group <group>
{forced-ref <configured frequency input timing
reference>} {threshold-ql <clock quality level>}
```

where

group <group>	is the frequency protection-group.
forced-ref <configured frequency input timing reference>	force switches to an input reference.
threshold-ql <clock quality level>	is the threshold quality level.

To set attributes for a phase protection-group

5 Set attributes for a phase protection-group:

```
sync phase protection-group set group <group> {forced-ref
<configured phase input timing reference>} {threshold-ql
<clock quality level>}
```

where

group <group>	is the phase protection-group.
forced-ref <configured phase input timing reference>	force switches to an input reference.
threshold-ql <clock quality level>	is the threshold quality level.

Copyright© 2022 Ciena® Corporation

To set attributes for a time-of-day protection-group

6 Set attributes for a time-of-day protection-group:

```
sync time-of-day protection-group set group <group>
{forced-ref <configured time-of-day input timing
reference>} {threshold-ql <clock quality level>}
```

where

group <group>	is the time-of-day protection-group
forced-ref <configured time-of-day input timing reference>	force switches to an input reference
threshold-ql <clock quality level>	is the threshold quality level

To add an input reference to a frequency protection-group

7 Add an input reference to a frequency protection-group:

```
sync frequency protection-group add group <group> {input-
ref <list of configured frequency input timing
references>}
```

where

group <group>	is the frequency protection-group
input-ref <list of configured frequency input timing references>	is the list of input references to add to the frequency protection-group

To add an input reference to phase protection-group

8 Add an input reference to a phase protection-group:

```
sync phase protection-group add group <group> input-ref
<list of configured phase input timing references>}
```

where

group <group>	is the phase protection-group
input-ref <list of configured phase input timing references>	is the list of input references to add to the phase protection-group

Copyright© 2022 Ciena® Corporation

To add an input reference to time-of-day protection-group

9 Add an input reference to a time-of-day protection-group:

```
sync time-of-day protection-group add group <group>
input-ref <list of configured time-of-day input timing
references>}
```

where

group <group>	is the time-of-day protection-group
input-ref <list of configured time-of-day input timing references>	is the list of input references to add to the time-of-day protection-group

To remove an input reference from frequency protection-group

10 Remove an input reference from a protection-group:

```
sync frequency protection-group remove group <group>
{input-ref <list of configured frequency input timing
references>}
```

where

group <group>	is the frequency protection-group
input-ref <list of configured frequency input timing references>	is the list of input references to remove from the frequency protection-group

To remove an input reference from a phase protection-group

11 Remove an input reference from a protection-group:

```
sync phase protection-group remove group <group> {input-
ref <list of configured phase input timing references>}
```

where

group <group>	is the phase protection-group
input-ref <list of configured phase input timing references>	is the list of input references to remove from the phase protection-group

Copyright© 2022 Ciena® Corporation

To remove an input reference from a time-of-day protection-group

12 Remove an input reference from a time-of-day protection-group:

```
sync time-of-day protection-group remove group <group>
{input-ref <list of configured time-of-day input timing
references>}
```

where

group <group>	is the time-of-day protection-group
input-ref <list of configured time-of-day input timing references>	is the list of input references to remove from the time-of-day protection-group

To clear a frequency protection-group's reference switch count

13 Clear a frequency protection-group's switch count:

```
sync frequency protection-group clear group <group-name>
ref-switch-count
```

where

group <group>	is the frequency protection-group

To clear a phase protection-group's reference switch count

14 Clear a phase protection-group's switch count:

```
sync phase protection-group clear group <group-name> ref-
switch-count
```

where

group <group>	is the phase protection-group

To clear a time-of-day protection-group's reference switch count

15 Clear a time-of-day protection-group's switch count:

```
sync time-of-day protection-group clear group <group-
name> ref-switch-count
```

where

group <group>	is the time-of-day protection-group

—end—

Copyright© 2022 Ciena® Corporation

Example

This example creates a frequency protection-group named FreqProtGroup.

```
sync frequency protection-group create group FreqProtGroup
threshold-ql sec
```

This example creates a time-of-day protection-group named TODGroup.

```
sync time-of-day protection-group create group TODGroup
threshold-ql sec
```

This example adds input references to the frequency protection-group named FreqProtGroup.

```
sync frequency protection-group add group FreqProtGroup input-
ref PTP_In,SyncE_Port5_In
```

This example adds input references to the time-of-day protection-group named TODGroup.

```
sync time-of-day protection-group add group TODGroup input-ref
myPTPinput5
```


Copyright© 2022 Ciena® Corporation

Procedure 5-17
Clearing timing statistics

You can clear

- SyncE timing statistics
- PTP master statistics
- PTP client statistics

Step	Action

To clear SyncE timing statistics

1 Clear SyncE timing statistics:

```
sync synce clear {statistics}
```

where

statistics — clears SyncE timing statistics.

To clear PTP master statistics

2 Clear PTP master statistics:

```
sync ptp input clear master <clock-ID> {pdu-statistics}
```

where

master <clock-ID> — is the clock ID for the PTP master.

pdu-statistics — clears PTP PDU statistics.

To clear PTP client statistics

3 Clear PTP client statistics:

```
sync ptp output clear client <clock-ID> {pdu-statistics}
```

where

client <clock-ID> — is the clock ID for the PTP client.

pdu-statistics — clears PTP PDU statistics.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 5-18
Displaying information for synchronization

You can display

- global synchronization information
- quality level hierarchy table, as defined in G.781
- all configured references

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 all configured references

3 Display all configured references:

```
sync show all
```

—end—

Example

This example shows sample output for the sync show command.

```
> sync show
+------------------------- SYNC CONFIGURATION ---------------------------+
| Parameter              | Value                                         |
+------------------------+-----------------------------------------------+
| Admin State            | Enabled                                       |
| Option Type            | Option II                                     |
| Reversion Mode         | Non-Revertive (default)                       |
| Wait-to-restore timer  | 5 minutes                                     |
| Holdover Interval      | Indefinite                                    |
+------------------------+-----------------------------------------------+

+--------------------- FREQUENCY OPERATIONAL STATE ------------------------+
| Parameter              | Value                                         |
+------------------------+-----------------------------------------------+
| Operational Status     | Locked                                        |
| Duration               | 1:01:10:10                                    |
| Operational QL         | ST3E  (last: ST2)                             |
| Selected Reference     | myBITSinput1                                  |
| Timing Interface       | sync-rj45-1                                   |
| Protocol               | BITS                                          |
| Ref Switch Count       | 2                                             |
+------------------------+-----------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

```
+------------------------- PHASE OPERATIONAL STATE -----------------------+
| Parameter              | Value                                          |
+------------------------+------------------------------------------------+
| Operational Status     | Locked                                         |
| Duration               | 0:02:12:40                                     |
| Operational QL         | ST3E (last: DNU)                               |
| Selected Reference     | myPTPinput5                                    |
| Timing Interface       | IP Interface     : PTP_Interface               |
|                        | Master IP Address: 10.10.10.1                  |
| Protocol               | PTP                                            |
| Ref Switch Count       | 1                                              |
+------------------------+------------------------------------------------+

+-------------------------TIME-OF-DAY-OPERATIONAL STATE ------------------+
| Parameter              | Value                                          |
+------------------------+------------------------------------------------+
| Operational Status     | Locked                                         |
| Duration               | 0:02:12:40                                     |
| Operational QL         | ST3E (last: DNU)                               |
| Selected Reference     | myPTPinput5                                    |
| Timing Interface       | IP Interface     : PTP_Interface               |
|                        | Master IP Address: 10.10.10.1                  |
| Protocol               | PTP                                            |
| Ref Switch Count       | 1                                              |
+------------------------+------------------------------------------------+
| Time-of-day            | Wed Dec 04 23:11:01 UTC 2013                   |
+------------------------+------------------------------------------------+
```

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

```
> sync show ql-hierarchy-table

+--- HIERARCHY OF QUALITY LEVELS ---+
| Option I  | Option II |   Order   |
+-----------+-----------+-----------+
| PRC       | PRS       |  Highest  |
| SSU-A     | STU       |     |     |
| SSU-B     | ST2       |     |     |
| SEC/EEC1 * | TNC       |     |     |
| DNU       | ST3E     *|     |     |
|           | ST3/EEC2  |     |     |
|           | SMC       |     |     |
|           | PROV      |     |     |
|           | DUS       |  Lowest   |
+-----------+-----------+-----------+
* Local OCXO QL
```


Copyright© 2022 Ciena® Corporation

This example shows sample output for the sync show all command.

```
> sync show all

+------------------------ FREQUENCY OPERATIONAL STATE -------------------------+
| Parameter              | Value                                                |
+------------------------+------------------------------------------------------+
| Operational Status     | Locked                                               |
| Duration               | 1:01:10:10                                           |
| Operational QL         | ST3E  (last: ST2)                                    |
| Selected Reference     | myBITSinput1                                         |
| Timing Interface       | sync-rj45-1                                          |
| Protocol               | BITS                                                 |
| Ref Switch Count       | 2                                                    |
+------------------------+------------------------------------------------------+


+------------------------------- FREQUENCY INPUT REFERENCES -------------------------------+
|Prot |                              |    Timing      |      |Over|Oper |    |                |
|Group|        Reference Name        |   Interface    |Proto|Prio| QL  |Prio| Operational Status |
+-----+------------------------------+----------------+-----+----+-----+----+--------------------+
| F-- |myBITSinput1                  |sync-rj45-1     |BITS |  1 |ST3E |  - |Active              |
| FPT |myPTPinput5                   |PTP_Interface   |PTP  |  - |ST3E |  3 |Active              |
| F-- |mySyncEinput5                 |port 5          |SyncE|  - |ST2  |  1 |Active              |
| F-- |mySyncEinput6                 |port 6          |SyncE|  - |ST2  |  2 |Active              |
| --- |myGPS10MHzInput1              |10mhz-smb-1     |GPS  |  - |PRS  |  - |Active              |
| --- |mySyncEinput8                 |port 8          |SyncE|  - |  -  |  4 |Not Authenticated   |
| --- |myTDMinput5                   |port tdm05      |TDM  |  - |  -  |  - |Loss of Signal      |
+-----+------------------------------+----------------+-----+----+-----+----+--------------------+

+-------------------------- FREQUENCY OUTPUT REFERENCES --------------------------+
|                              |    Timing      |      |Oper |                    |
|        Reference Name        |   Interface    |Proto| QL  | Operational Status |
+------------------------------+----------------+-----+-----+--------------------+
|myGPS10MHzOutput2             |10mhz-smb-1     |GPS  |  -  |Active              |
|myBITSoutput2                 |sync-rj45-1     |BITS |ST3E |Active              |
|myPTPoutput9                  |PTP_Interfac...|PTP  |ST3E |Active              |
|mySyncEoutput1                |port 1          |SyncE|  -  |Loss of Signal      |
+------------------------------+----------------+-----+-----+--------------------+

+-------------------------- PHASE OPERATIONAL STATE -------------------------+
| Parameter              | Value                                                |
+------------------------+------------------------------------------------------+
| Operational Status     | Locked                                               |
| Duration               | 0:02:12:40                                           |
| Operational QL         | ST3E  (last: DNU)                                    |
| Selected Reference     | myPTPinput5                                          |
| Timing Interface       | IP Interface      : PTP_Interface                    |
|                        | Master IP Address : 10.10.10.1                       |
| Protocol               | PTP                                                  |
| Ref Switch Count       | 1                                                    |
+------------------------+------------------------------------------------------+

+---------------------------------- PHASE INPUT REFERENCES ----------------------------------+
|Prot |                              |    Timing      |      |Over|Oper |    |                |
|Group|        Reference Name        |   Interface    |Proto|Prio| QL  |Prio| Operational Status |
+-----+------------------------------+----------------+-----+----+-----+----+--------------------+
| FPT |myPTPinput5                   |PTP_Interface   |PTP  |  - |ST3E |  3 |Active              |
+-----+------------------------------+----------------+-----+----+-----+----+--------------------+

+---------------------------- PHASE OUTPUT REFERENCES ----------------------------+
|                              |    Timing      |      |Oper |                    |
|        Reference Name        |   Interface    |Proto| QL  | Operational Status |
+------------------------------+----------------+-----+-----+--------------------+
|myPTPoutput9                  |PTP_Interfac...|PTP  |ST3E |Active              |
+------------------------------+----------------+-----+-----+--------------------+
```


Copyright© 2022 Ciena® Corporation

```
+------------------------ TIME-OF-DAY OPERATIONAL STATE ------------------------+
| Parameter              | Value                                                 |
+------------------------+-------------------------------------------------------+
| Operational Status     | Locked                                                |
| Duration               | 0:02:12:40                                            |
| Operational QL         | ST3E  (last: DNU)                                     |
| Selected Reference     | myPTPinput5                                           |
| Timing Interface       | IP Interface      : PTP_Interface                     |
|                        | Master IP Address : 10.10.10.1                        |
| Protocol               | PTP                                                   |
| Ref Switch Count       | 1                                                     |
+------------------------+-------------------------------------------------------+
| Time-of-day            | Wed Dec 04 23:11:01 UTC 2013                          |
+------------------------+-------------------------------------------------------+


+--------------------------------- TIME-OF-DAY INPUT REFERENCES -----------------------------+
|Prot |                              |    Timing     |      |Over|Oper |    |                    |
|Group|        Reference Name        |   Interface   |Proto|Prio| QL  |Prio| Operational Status |
+-----+------------------------------+---------------+-----+----+-----+----+--------------------+
| FPT |myPTPinput5                   |PTP_Interface  |PTP  |  - |ST3E |  3 |Active              |
+-----+------------------------------+---------------+-----+----+-----+----+--------------------+

+------------------------- TIME-OF-DAY OUTPUT REFERENCES -------------------------+
|                                |    Timing     |      |Oper |                    |
|        Reference Name          |   Interface   |Proto| QL  | Operational Status |
+--------------------------------+---------------+-----+-----+--------------------+
|myPTPoutput9                    |PTP_Interfac...|PTP  |ST3E |Active              |
+--------------------------------+---------------+-----+-----+--------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 5-19
Displaying SyncE information

You can display:

- summary information for SyncE
- statistics information for SyncE
- information for SyncE input references
- detailed information for a SyncE input reference
- information for SyncE output references
- detailed information for a SyncE output reference

Step	Action

To display summary information for SyncE

1 Display summary information for synchronous Ethernet:

```
sync synce show
```

To display statistics information for SyncE

2 Display statistics information for synchronous Ethernet:

```
sync synce show statistics
```

To display information for SyncE input references

3 Display information for synchronous Ethernet input references:

```
sync synce input show
```

To display detailed information for a SyncE input reference

4 Display detailed information for a SyncE input reference:

```
sync synce input show ref <ref-name>
```

where

ref <ref-name> is the SyncE input reference to display detailed information for.

To display information for SyncE output references

5 Display information for SyncE output references:

```
sync synce output show
```


Copyright© 2022 Ciena® Corporation

To display detailed information for a SyncE output reference

6 Display detailed information for a SyncE output reference:

```
sync synce output show ref <ref-name>
```

where

ref <ref-name> is the SyncE output reference to display detailed information for.

—end—

Example

This example shows sample output for the sync synce output show command.

```
> sync synce output show

+-------------------------- SYNCE OUTPUT REFERENCES ----------------------------+
|                              |    Timing     |      |Oper |                    |
|        Reference Name        |   Interface   |Proto|QL/CC| Operational Status |
+------------------------------+---------------+-----+-----+--------------------+
|SYNCE_OUT_3                   |port 3         |SyncE|PRC  |Active              |
+------------------------------+---------------+-----+-----+--------------------+
```

This example shows sample output for the sync synce output show ref command for the output reference SYNCE_OUT_3.

```
> sync synce output show ref SYNCE_OUT_3

+-------------------------- SYNCE OUTPUT REFERENCE ----------------------------+
| Parameter              | Value                                               |
+------------------------+-----------------------------------------------------+
| Reference Name         | SYNCE_OUT_3                                         |
| Timing Interface       | port 3                                              |
| Operational Status     | Active                                              |
+------------------------+-----------------------------------------------------+
| Protocol               | SyncE                                               |
+------------------------+-----------------------------------------------------+
| TX QL                  | PRC                                                 |
+------------------------+-----------------------------------------------------+
```

This example shows sample output for the sync synce show statistics command.

```
> sync synce show statistics

+-------+------------ SYNCE STATISTICS ---------------------+
|       |             Rx            |             Tx          |
|  QL   +-------------+-------------+-------------+-------------+
|       |   Info Msg  |  Event Msg  |   Info Msg  |  Event Msg  |
+-------+-------------+-------------+-------------+-------------+
| PRS   |           0|           0|           0|           0|
| STU   |           0|           0|           0|           0|
| ST2   |           0|           0|           0|           0|
| TNC   |           0|           0|           0|           0|
| ST3E  |           0|           0|           0|           0|
| ST3   |           0|           0|           0|           0|
| SMC   |           0|           0|           0|           0|
| PROV  |           0|           0|           0|           0|
| DUS   |           0|           0|           0|           0|
+-------+-------------+-------------+-------------+-------------+
```


Copyright© 2022 Ciena® Corporation

This example shows sample output for the sync synce input show command.

```
> sync synce input show

+----------------------------- SYNCE INPUT REFERENCES ----------------------------+
|Prot |                      |  Timing    |      |Over|Oper |    |                    |
|Group|  Reference Name      | Interface  |Proto|Prio| QL  |Prio| Operational Status |
+-----+----------------------+------------+-----+----+-----+----+--------------------+
| F-- |mySyncEinput5         |port 5      |SyncE|  - |ST2  |  1 |Active              |
| F-- |mySyncEinput6         |port 6      |SyncE|  - |ST2  |  2 |Active              |
| --- |mySyncEinput8         |port 8      |SyncE|  - |  -  |  4 |Not Authenticated   |
+-----+----------------------+------------+-----+----+-----+----+--------------------+
```

This example shows sample output for the sync synce input show ref command for the input reference mySyncEinput5.

```
> sync synce input show ref mySyncEinput5

+------------------------- SYNCE INPUT REFERENCE ------------------------+
| Parameter                | Value                                       |
+-------------------------+---------------------------------------------+
| Reference Name           | mySyncEinput5                               |
| Timing Interface         | port 5                                      |
| Protection Group         | myFreqGroup                                 |
| Operational Status       | Active                                      |
+-------------------------+---------------------------------------------+
| Override-priority        | -                                           |
| Operational QL           | ST2                                         |
| Forced QL                | -                                           |
| RX QL                    | ST2                                         |
| QL Receive               | Enabled                                     |
| Priority                 | 1                                           |
+-------------------------+---------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 5-20
Displaying BITS information

You can display:

- summary information for BITS timing
- information for BITS input references
- detailed information for a BITS input reference
- information for BITS output references
- detailed information for a BITS output reference

Step	Action

To display summary information for BITS timing

1 Display summary information for BITS timing:

```
sync bits show
```

To display information for BITS input references

2 Display information for BITS input references:

```
sync bits input show
```

To display detailed information for a BITS input reference

3 Display detailed information for a BITS input reference:

```
sync bits input show ref <ref>
```

To display information for BITS output references

4 Display information for BITS output references:

```
sync bits output show
```

To display detailed information for a BITS output reference

5 Display detailed information for a BITS output reference:

```
sync bits output show ref <ref>
```

—end—

Copyright© 2022 Ciena® Corporation

Example

This example shows sample output for the sync bits show command.

```
> sync bits show

+------------------------- BITS INPUT REFERENCES -------------------------------+
|Prot |                   |     Timing      |      |Over|Oper |    |                |
|Group|  Reference Name   |   Interface     |Proto|Prio| QL  |Prio| Operational Status |
+-----+-----------------+---------------+-----+----+-----+----+-------------------+
| F-- |myBITSinput1      |sync-rj45-1    |BITS |  1 |ST3E |  - |Active             |
+-----+-----------------+---------------+-----+----+-----+----+-------------------+


+-------------------------- BITS OUTPUT REFERENCES ------------------------+
|                           |     Timing      |      |Oper |                |
|     Reference Name        |   Interface     |Proto| QL  | Operational Status |
+---------------------------+---------------+-----+-----+-------------------+
|myBITSoutput2              |sync-rj45-2    |BITS |ST3E |Active             |
+---------------------------+---------------+-----+-----+-------------------+
```

This example shows sample output for the sync bits input show command.

```
> sync bits input show

+-------------------------- BITS INPUT REFERENCES -------------------------------+
|Prot |                   |     Timing      |      |Over|Oper |    |                |
|Group|  Reference Name   |   Interface     |Proto|Prio| QL  |Prio| Operational Status |
+-----+-----------------+---------------+-----+----+-----+----+-------------------+
| F-- |myBITSinput1      |sync-rj45-1    |BITS |  1 |ST3E |  - |Active             |
+-----+-----------------+---------------+-----+----+-----+----+-------------------+
```

This example shows sample output for the sync bits input show ref command for the input reference myBITSinput1.

```
> sync bits input show ref myBITSinput1

+------------------------ BITS INPUT REFERENCE ------------------------+
| Parameter                | Value                                      |
+-------------------------+--------------------------------------------+
| Reference Name           | myBITSinput1                               |
| Timing Interface         | sync-rj45-1                                |
| Protection Group         | myFreqGroup                                |
| Operational Status       | Active                                     |
+-------------------------+--------------------------------------------+
| BITS Mode                | E1                                         |
| Format                   | ESF                                        |
| Encoding                 | AMI                                        |
| E1 SSM Location          | SA4                                        |
+-------------------------+--------------------------------------------+
| Override Priority        | 1                                          |
| Operational QL           | ST3E                                       |
| Forced QL                | ST3E                                       |
| RX QL                    | -                                          |
| QL Receive               | Disabled                                   |
| Priority                 | -                                          |
+-------------------------+--------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

This example shows sample output for the sync bits output show command.

```
> sync bits output show
+------------------------- BITS OUTPUT REFERENCES ------------------------+
|                          |    Timing     |      |Oper |                  |
|        Reference Name    |   Interface   |Proto| QL  | Operational Status |
+--------------------------+---------------+-----+-----+-------------------+
|myBITSoutput2             |sync-rj45-2    |BITS |ST3E |Active             |
+--------------------------+---------------+-----+-----+-------------------+
```

This example shows sample output for the sync bits output show ref command for the output reference myBITSinput2.

```
> sync bits output show ref myBITSinput2

+------------------------- BITS OUTPUT REFERENCE ------------------------+
| Parameter              | Value                                         |
+------------------------+-----------------------------------------------+
| Reference Name         | myBITSinput2                                  |
| Timing Interface       | sync-rj45-2                                   |
| Operational Status     | Active                                        |
+------------------------+-----------------------------------------------+
| BITS Mode              | E1                                            |
| Format                 | ESF                                           |
| Encoding               | AMI                                           |
| E1 SSM Location        | SA4                                           |
+------------------------+-----------------------------------------------+
| TX QL                  | ST3E                                          |
+------------------------+-----------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 5-21
Displaying PTP information

You can display:

- summary information for PTP timing
- performance statistics information for PTP timing
- information for PTP input references
- information for PTP output references
- detailed information for a PTP input reference
- detailed information for a PTP output reference
- information for PTP masters
- information for PTP clients
- PDU statistics for a PTP master
- PDU statistics for a PTP client

Step	Action

To display summary information for PTP timing

1 Display summary information for PTP timing:

```
sync ptp show
```

To display performance statistics information for PTP timing

2 Display performance statistics information for PTP timing:

```
sync ptp show performance-statistics
```

To display information for PTP input references

3 Display information for PTP input references:

```
sync ptp input show
```

To display detailed information for a PTP input reference

4 Display detailed information for a PTP input reference:

```
sync ptp input show ref <ref-name>
```

To display the list of connected PTP masters

5 Display the list of connected PTP masters:

```
sync ptp input show master-list
```

To display the PDU statistics for a PTP master

6 Display the PDU statistics for a PTP master:

```
sync ptp input show master <clock-ID> pdu-statistics
```


Copyright© 2022 Ciena® Corporation

To display information for PTP output references

7 Display information for PTP output references:

```
sync ptp output show
```

To display detailed information for a PTP output reference

8 Display detailed information for a PTP output reference:

```
sync ptp output show ref <ref-name>
```

To display the connected PTP client list

9 Display the connected PTP client list:

```
sync ptp output show client-list
```

To display the PDU statistics for a PTP client

10 Display the PDU statistics for a PTP client:

```
sync ptp output show client <clock-ID> pdu-statistics
```

—end—

Example

This example shows sample output for the sync ptp show command with profile as `G.8265.1`.

```
> sync ptp show

+---------------------------- PTP CONFIGURATION ------------------------+
| Parameter                | Value                                       |
+-------------------------+---------------------------------------------+
| Clock Type               | Boundary Clock                              |
| Clock Identity           | 2c:39:c1:ff:fe:cc:ec:60                     |
| Address Mode             | Unicast                                     |
| Tag Priority             | 7                                           |
| Protocol Version         | 2                                           |
| Profile                  | G.8265.1                                    |
| Profile Version          | 1.0                                         |
| Profile Identifier       | 00:19:a7:00:01:00                           |
| Domain Number            | 4                                           |
| DSCP                     | 56                                          |
+-------------------------+---------------------------------------------+
+--------------------------- PTP INPUT CONFIGURATION ----------------------+
| Parameter                | Value                                       |
+-------------------------+---------------------------------------------+
| Unicast Req Duration     | 300                                         |
| Announce Rx Timeout      | 3                                           |
| Announce Tx Rate Req     | 1-2nd                                       |
| Sync Tx Rate Req         | 64 pps                                      |
| Delay Req Tx Rate        | 64 pps                                      |
+-------------------------+---------------------------------------------+

+------------------------------- PTP INPUT REFERENCES ---------------------------+
|Prot |                |     Timing       |      |Over|Oper |    |                    |
|Group| Reference Name |    Interface     |Proto|Prio| QL  |Prio| Operational Status |
+-----+----------------+-----------------+-----+----+-----+----+--------------------+
|FPT- |myPTPinput5     |PTP_Interface    |PTP  |  - |ST3E |  3 |Active              |
+-----+----------------+-----------------+-----+----+-----+----+--------------------+
```


Copyright© 2022 Ciena® Corporation

```
+-------------------------- PTP OUTPUT CONFIGURATION --------------------+
| Parameter              | Value                                          |
+------------------------+------------------------------------------------+
| Max Slave Sessions     | 16                                             |
| Timestamp Mode         | Two-step                                       |
+------------------------+------------------------------------------------+

+---------------------------- PTP OUTPUT REFERENCES --------------+
|                 |     Timing      |      |Over |                  |
| Reference Name  |   Interface     |Proto| QL  | Operational Status |
+----------------+----------------+-----+-----+--------------------+
| myPTPinput9     |PTP Interfac...|PTP  |ST3E | Active             |
+----------------+----------------+-----+-----+--------------------+
```

This example shows sample output for the sync ptp show command with profile as `G.8275.1`.

```
              sync ptp show
+----------------------------- PTP CONFIGURATION -----------------------+
| Parameter              | Value                                          |
+------------------------+------------------------------------------------+
| Clock Type             | Boundary Clock                                 |
| Clock Identity         | 54:c3:3e:ff:fe:4f:12:c0                        |
| Address Mode           | Multicast                                      |
| Protocol Version       | 2                                              |
| Profile                | G.8275.1                                       |
| Profile Version        | 1.0                                            |
| Profile Identifier     | 00:19:a7:01:01:00                              |
| Domain Number          | 24                                             |
| Destination Mac        | 01:1b:19:00:00:00 (Forwardable)                |
+------------------------+------------------------------------------------+

+-------------------------- PTP INPUT CONFIGURATION --------------------+
| Parameter              | Value                                          |
+------------------------+------------------------------------------------+
| Announce Rx Loss Num   | 3 announce intervals                           |
| Delay Req Rate         | 16 pps                                         |
+------------------------+------------------------------------------------+

+---------------------------- PTP INPUT REFERENCES ----------------------------------+
|Prot |                        |    Timing      |      |Over|Oper |    |                |
|Group|        Reference Name  |   Interface    |Proto|Prio|QL/CC|Prio| Operational Status |
+-----+------------------------+----------------+-----+----+-----+----+--------------------+
| -P- |ptp_in_1                |port 1          |PTP  |-   |6    |-   |Active              |
+-----+------------------------+----------------+-----+----+-----+----+--------------------+

+-------------------------- PTP OUTPUT CONFIGURATION --------------------+
| Parameter              | Value                                          |
+------------------------+------------------------------------------------+
| Timestamp Mode         | One-Step                                       |
| Priority1              | 128                                            |
| Priority2              | 128                                            |
| Announce Rate          | 8 pps                                          |
| Sync Rate              | 16 pps                                         |
+------------------------+------------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

```
+---------------------------- PTP OUTPUT REFERENCES ----------------------+
|                          |    Timing    |     |Oper |                    |
|        Reference Name    |   Interface  |Proto|QL/CC| Operational Status |
+-------------------------+--------------+-----+-----+--------------------+
|ptp_out_2                |port 2        |PTP  |248  |Active              |
+-----------------------------+---------------+-----+-----+-------------+
```

This command shows sample output for the sync ptp show performance-statistics command.

```
> sync ptp show performance-statistics

+------------------ PTP PERFORMANCE STATISTICS -------------------+
| Statistic                                 | Value                |
+-------------------------------------------+----------------------+
| Frequency correction (ppb)                | 20.471508            |
| Phase correction (ns)                     | 0.000000             |
| Minimal round trip delay (ns)             | 0.000000             |
| Fwd Flow operational min TDEV (ns)        | 3725.350342          |
| Rev flow operational min TDEV (ns)        | 2298.259033          |
+-------------------------------------------+----------------------+
```

This command shows sample output for the sync ptp input show command.

```
> sync ptp input show

+--------------------------- PTP INPUT CONFIGURATION --------------------------+
| Parameter              | Value                                                |
+------------------------+------------------------------------------------------+
| Unicast Req Duration   | 300                                                  |
| Announce Rx Timeout    | 3 sec                                                |
| Announce Rx Rate Req   | 1-2nd                                                |
| Sync Tx Rate Req       | 64 pps                                               |
| Delay Req Tx Rate      | 64 pps                                               |
+------------------------+------------------------------------------------------+

+------------------------------------ PTP INPUT REFERENCES ----------------------------------+
|Prot |                              |    Timing     |      |Over|Oper |    |                  |
|Group|        Reference Name        |   Interface   |Proto|Prio| QL  |Prio| Operational Status |
+-----+------------------------------+---------------+-----+----+-----+----+-------------------+
| FPT |myPTPinput5                   |PTP_Interface  |PTP  |  - |ST3E |  3 |Active             |
+-----+------------------------------+---------------+-----+----+-----+----+-------------------+
```


Copyright© 2022 Ciena® Corporation

This command shows sample output for the sync ptp input show ref myPTPinput5 command.

```
> sync ptp input show ref myPTPinput5

+---------------------------- PTP INPUT REFERENCE ----------------------------+
| Parameter              | Value                                               |
+------------------------+-----------------------------------------------------+
| Reference Name         | myPTPinput5                                         |
| Timing Interface       | IP Interface      : PTP_Interface                   |
|                        | Master IP Address : 10.10.10.1                      |
| Protection Group       | myFreqGroup, myPhaseGroup, myTodGroup               |
| Operational Status     | Active                                              |
+------------------------+-----------------------------------------------------+
| Encapsulation Type     | UDP-over-IP                                         |
+------------------------+-----------------------------------------------------+
| Override Priority      | -                                                   |
| Oper Clock Class/QL    | 100/ST3E                                            |
| Forced Clock Class/QL  | -                                                   |
| RX Clock Class/QL      | 100/ST3E                                            |
| QL Receive             | Enabled                                             |
| Priority               | 3                                                   |
+------------------------+-----------------------------------------------------+
```

This command shows the sample output for the sync ptp input show master-list command.

```
> sync ptp input show master-list

+--------------------------------- PTP MASTER LIST ---------------------------------+
|                        | Announce| Announce|   Sync  |   Sync  |Delay Res|Delay Res|
|    Clock Identity      |Rate(pps)|Dur (sec)|Rate(pps)|Dur (sec)|Rate(pps)|Dur (sec)|
+------------------------+---------+---------+---------+---------+---------+--------+
|00:11:22:33:44:55:66:77|       64|      300|       64|      300|       64|      300|
+------------------------+---------+---------+---------+---------+---------+--------+
```


Copyright© 2022 Ciena® Corporation

This command shows sample output for the sync ptp input show master 00.11.22.33.44.55.66.77 pdu-statistics command.

```
> sync ptp input show master 00:11:22:33:44:55:66:77 pdu-statistics

+------------------------------ PTP MASTER ENTRY --------------------------+
| Parameter              | Value                                            |
+------------------------+--------------------------------------------------+
| Clock Identity         | 00:11:22:33:44:55:66:77                          |
| IP Address             | 10.10.10.1                                       |
+------------------------+--------------------------------------------------+

+---------------- PTP MASTER PDU STATISTICS ----------------+
| Statistic                          | Value                 |
+------------------------------------+-----------------------+
| Rx Announce                        | 0                     |
| Rx Sync                            | 0                     |
| Rx Follow-up                       | 0                     |
| Rx Delay Request                   | 0                     |
| Rx Delay Response                  | 0                     |
| Rx Signalling                      | 0                     |
| Rx Management                      | 0                     |
| Tx Announce                        | 0                     |
| Tx Sync                            | 0                     |
| Tx Follow-up                       | 0                     |
| Tx Delay Request                   | 0                     |
| Tx Delay Response                  | 0                     |
| Tx Signalling                      | 0                     |
| Tx Management                      | 0                     |
+------------------------------------+-----------------------+
```

This command shows sample output for the sync ptp output show command.

```
> sync ptp output show

+------------------------- PTP OUTPUT CONFIGURATION -------------------+
| Parameter              | Value                                      |
+------------------------+--------------------------------------------+
| Max Slave Sessions     | 16                                         |
+------------------------+--------------------------------------------+
+----------------------------- PTP OUTPUT REFERENCES ----------------------------+
|                              |    Timing      |      |Oper |                    |
|        Reference Name        |   Interface    |Proto| QL  | Operational Status |
+------------------------------+----------------+-----+-----+--------------------+
|myPTPoutput9                  |PTP_Interfac...|PTP  |ST3E |Active              |
+------------------------------+----------------+-----+-----+--------------------+
```


Copyright© 2022 Ciena® Corporation

This command shows sample output for the sync ptp output show ref myPTPoutput5 command.

```
> sync ptp output show ref myPTPoutput5

+--------------------------- PTP OUTPUT REFERENCE --------------------------+
| Parameter                | Value                                          |
+------------------------+-------------------------------------------------+
| Reference Name           | myPTPoutput5                                   |
| Timing Interface         | IP Interface     : PTP_InterfaceOut            |
| Operational Status       | Active                                         |
+------------------------+-------------------------------------------------+
| Encapsulation Type       | UDP-over-IP                                    |
+------------------------+-------------------------------------------------+
| TX Clock Class/QL        | 100/ST3E                                       |
+------------------------+-------------------------------------------------+
```

This command shows sample output for the sync ptp-output show client-list command:

```
> sync ptp output show client-list

+------------------------------- PTP CLIENT LIST --------------------------------+
|                       | Announce| Announce|    Sync   |    Sync   |Delay Res|Delay Res|
|    Clock Identity     |Rate(pps)|Dur (sec)|Rate(pps)|Dur (sec)|Rate(pps)|Dur (sec)|
+-----------------------+---------+---------+---------+---------+---------+---------+
|11:22:33:44:55:66:77:88|       64|      300|       64|      300|       64|      300|
|22:33:44:55:66:77:88:99|       64|      300|       64|      300|       64|      300|
|33:44:55:66:77:88:99:AA|       64|      300|       64|      300|       64|      300|
+-----------------------+---------+---------+---------+---------+---------+---------+
```

This command shows a sample output for the sync ptp output show client 11:22:33:44:55:66:77:88 pdu-statistics command:

```
> sync ptp output show client 11:22:33:44:55:66:77:88 pdu-statistics

+----------------------------- PTP CLIENT ENTRY ------------------------------+
| Parameter              | Value                                              |
+-----------------------+-------------------------------------------------+
| Clock Identity         | 11:22:33:44:55:66:77:88                            |
| IP Address             | 10.10.10.5                                         |
+-----------------------+-------------------------------------------------+
+---------------- PTP CLIENT PDU STATISTICS ----------------+
| Statistic                          | Value                 |
+-----------------------------------+-----------------------+
| Rx Announce                        | 0                     |
| Rx Sync                            | 0                     |
| Rx Follow-up                       | 0                     |
| Rx Delay Request                   | 0                     |
| Rx Delay Response                  | 0                     |
| Rx Signalling                      | 0                     |
| Rx Management                      | 0                     |
| Tx Announce                        | 0                     |
| Tx Sync                            | 0                     |
| Tx Follow-up                       | 0                     |
| Tx Delay Request                   | 0                     |
| Tx Delay Response                  | 0                     |
| Tx Signalling                      | 0                     |
| Tx Management                      | 0                     |
+-----------------------------------+-----------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 5-22
Displaying GPS information

You can display:

- summary information for GPS timing
- information for GPS input references
- detailed information for a GPS input reference
- information for GPS output references
- detailed information for a GPS output reference

Step	Action

To display summary information for GPS timing

1 Display summary information for GPS timing:

```
sync gps show
```

To display information for GPS input references

2 Display information for GPS input references:

```
sync gps input show
```

To display detailed information for a GPS input reference

3 Display detailed information for a GPS input reference:

```
sync gps input show ref <ref-name>
```

To display information for GPS output references

4 Display information for GPS output references:

```
sync gps output show
```

To display detailed information for a GPS output reference

5 Display detailed information for a GPS output reference:

```
sync gps output show ref <ref-name>
```

—end—

Copyright© 2022 Ciena® Corporation

Example

This command shows sample output for the sync gps show command.

```
> sync gps show

+------------------------------------ GPS INPUT REFERENCES ------------------------------------+
|Prot |                                |    Timing      |      |Over|Oper |    |                  |
|Group|        Reference Name          |   Interface    |Proto|Prio| QL  |Prio| Operational Status |
+-----+--------------------------------+----------------+-----+----+-----+----+-------------------+
| --- |myGPS10MHzInput1                |10mhz-smb-1     |GPS  |  - |PRS  |  - |Active             |
+-----+--------------------------------+----------------+-----+----+-----+----+-------------------+

+--------------------------- GPS OUTPUT CONFIGURATION --------------------------+
| Parameter              | Value                                                 |
+------------------------+-------------------------------------------------------+
| 1pps Pulse Width       | 1000 usec                                             |
+------------------------+-------------------------------------------------------+

+----------------------------- GPS OUTPUT REFERENCES -----------------------------+
|                                |    Timing      |      |Oper |                   |
|        Reference Name          |   Interface    |Proto| QL  | Operational Status |
+--------------------------------+----------------+-----+-----+-------------------+
|myGPS10MHzOutput2               |10mhz-smb-1     |GPS  |  -  |Active             |
+--------------------------------+----------------+-----+-----+-------------------+
```

This command shows sample output for the sync gps input show command.

```
> sync gps input show

+------------------------------------ GPS INPUT REFERENCES ------------------------------------+
|Prot |                                |    Timing      |      |Over|Oper |    |                  |
|Group|        Reference Name          |   Interface    |Proto|Prio| QL  |Prio| Operational Status |
+-----+--------------------------------+----------------+-----+----+-----+----+-------------------+
| --- |myGPS10MHzInput1                |10mhz-smb-1     |GPS  |  - |PRS  |  - |Active             |
+-----+--------------------------------+----------------+-----+----+-----+----+-------------------+
```

This command shows sample output for the sync gps input show ref myGPS10MHzInput1 command.

```
> sync gps input show ref myGPS10MHzInput1

+--------------------------- GPS INPUT REFERENCE ---------------------------+
| Parameter              | Value                                           |
+------------------------+-------------------------------------------------+
| Reference Name         | myGPS10MHzInput1                                |
| Timing Interface       | 10mhz-smb-1                                     |
| Protection Group       | -                                               |
| Operational Status     | Active                                          |
+------------------------+-------------------------------------------------+
| Clock Mode             | Frequency                                       |
| Frequency Clock        | 10 MHz                                          |
+------------------------+-------------------------------------------------+
| Override Priority      | -                                               |
| Operational QL         | PRS                                             |
| Forced QL              | PRS                                             |
| Priority               | -                                               |
+------------------------+-------------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

This example shows sample output for the sync gps output show command.

```
> sync gps output show

+------------------------- GPS OUTPUT CONFIGURATION -------------------------+
| Parameter              | Value                                             |
+------------------------+---------------------------------------------------+
| 1pps Pulse Width       | 1000 usec                                         |
+------------------------+---------------------------------------------------+

+---------------------------- GPS OUTPUT REFERENCES ----------------------------+
|                                |     Timing      |      |Oper |                 |
|        Reference Name          |   Interface     |Proto| QL  | Operational Status |
+--------------------------------+-----------------+-----+-----+-------------------+
|myGPS10MHzOutput2               |10mhz-smb-1      |GPS  |  -  |Active             |
+--------------------------------+-----------------+-----+-----+-------------------+
```

This example shows sample output for the sync gps output show ref myGPS10MHzOutput2 command.

```
> sync gps output show ref myGPS10MHzOutput2

+--------------------------- GPS OUTPUT REFERENCE -------------------------+
| Parameter              | Value                                            |
+------------------------+--------------------------------------------------+
| Reference Name         | myGPS10MHzOutput2                                |
| Timing Interface       | 10mhz-smb-1                                      |
| Operational Status     | Active                                           |
+------------------------+--------------------------------------------------+
| Clock Mode             | Frequency                                        |
| Frequency Clock        | 10 MHz                                           |
+------------------------+--------------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 5-23
Displaying TDM information

You can display:

- summary information for TDM timing
- summary information for TDM input references
- detailed information for a TDM input reference

Step	Action

To display summary information for TDM timing

1 Display TDM information:

```
sync tdm show
```

To display summary information for TDM input references

2 Display summary information for TDM input references:

```
sync tdm input show
```

To display detailed information for a TDM input reference

3 Display detailed information for a TDM input reference:

```
sync tdm input show ref <ref-name>
```

—end—

Copyright© 2022 Ciena® Corporation

Example

This example shows sample output for the sync tdm show command.

```
> sync tdm show

+----------------------------- TDM INPUT REFERENCES -----------------------------+
|Prot |                  |     Timing      |      |Over|Oper |    |                    |
|Group| Reference Name   |   Interface     |Proto|Prio| QL  |Prio| Operational Status |
+-----+------------------+-----------------+-----+----+-----+----+--------------------+
| --- |myTDMinput5       |port tdm05       |TDM  |  - |  -  |  - |Loss of Signal      |
+-----+------------------+-----------------+-----+----+-----+----+--------------------+
```

This example shows sample output for the sync tdm input show command.

```
> sync tdm input show

+----------------------------- TDM INPUT REFERENCES -----------------------------+
|Prot |                  |     Timing      |      |Over|Oper |    |                    |
|Group| Reference Name   |   Interface     |Proto|Prio| QL  |Prio| Operational Status |
+-----+------------------+-----------------+-----+----+-----+----+--------------------+
| --- |myTDMinput5       |port tdm05       |TDM  |  - |  -  |  - |Loss of Signal      |
+-----+------------------+-----------------+-----+----+-----+----+--------------------+
```

This example shows sample output for the sync tdm input show ref myTDMinput5 command.

```
> sync tdm input show ref myTDMinput5

+-------------------------- TDM INPUT REFERENCE --------------------------+
| Parameter                | Value                                        |
+------------------------+-----------------------------------------------+
| Reference Name           | myTDMinput5                                  |
| Timing Interface         | port 5                                       |
| Protection Group         | -                                            |
| Operational Status       | Loss of Signal                               |
+------------------------+-----------------------------------------------+
| Override-priority        | -                                            |
| Operational QL           | -                                            |
| Forced QL                | ST2                                          |
| Priority                 | -                                            |
+------------------------+-----------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 5-24
Displaying frequency information

You can display

- frequency information
- summary information for frequency input references
- summary information for frequency output references

Step	Action

To display detailed information for frequency

1 Display detailed information for frequency:

```
sync frequency show
```

To display summary information for frequency input references

2 Display summary information for frequency input references:

```
sync frequency input show
```

To display summary information for frequency output references

3 Display summary information for output references for frequency:

```
sync frequency output show
```

—end—

Example

This example shows sample output for the sync frequency show command.

```
> sync frequency show

+------------------------ FREQUENCY OPERATIONAL STATE -------------------------+
| Parameter              | Value                                                |
+------------------------+------------------------------------------------------+
| Operational Status     | Locked                                               |
| Duration               | 1:01:10:10                                           |
| Operational QL         | ST3E  (last: ST2)                                    |
| Selected Reference     | myBITSinput1                                         |
| Timing Interface       | sync-rj45-1                                          |
| Protocol               | BITS                                                 |
| Ref Switch Count       | 2                                                    |
+------------------------+------------------------------------------------------+
+---------------------------------- FREQUENCY INPUT REFERENCES ----------------------------------+
|Prot |                                |    Timing      |      |Over|Oper |    |                    |
|Group|        Reference Name          |   Interface    |Proto|Prio| QL  |Prio| Operational Status |
+-----+--------------------------------+----------------+-----+----+-----+----+--------------------+
| F-- |myBITSinput1                    |sync-rj45-1     |BITS |  1 |ST3E |  - |Active              |
| FPT |myPTPinput5                     |PTP_Interface   |PTP  |  - |ST3E |  3 |Active              |
| F-- |mySyncEinput5                   |port 5          |SyncE|  - |ST2  |  1 |Active              |
| F-- |mySyncEinput6                   |port 6          |SyncE|  - |ST2  |  2 |Active              |
| --- |myGPS10MHzInput1                |10mhz-smb-1     |GPS  |  - |PRS  |  - |Active              |
| --- |mySyncEinput8                   |port 8          |SyncE|  - |  -  |  4 |Not Authenticated   |
| --- |myTDMinput5                     |port tdm05      |TDM  |  - |  -  |  - |Loss of Signal      |
+-----+--------------------------------+----------------+-----+----+-----+----+--------------------+
```


Copyright© 2022 Ciena® Corporation

```
+------------------------- FREQUENCY OUTPUT REFERENCES ------------------------+
|                              |     Timing     |       |Oper |                    |
|        Reference Name        |   Interface    |Proto  | QL  | Operational Status |
+------------------------------+----------------+-----+-----+--------------------+
|myGPS10MHzOutput2             |10mhz-smb-1     |GPS  |  -  |Active              |
|myBITSoutput2                 |sync-rj45-1     |BITS |ST3E |Active              |
|myPTPoutput9                  |PTP_Interfac...|PTP  |ST3E |Active              |
|mySyncEoutput1                |port 1          |SyncE|  -  |Loss of Signal      |
+------------------------------+----------------+-----+-----+--------------------+
```

This example shows sample output for the sync frequency input show command.

```
> sync frequency input show
```

```
+--------------------------------- FREQUENCY INPUT REFERENCES ---------------------------------+
|Prot |                              |     Timing     |       |Over|Oper |    |                    |
|Group|        Reference Name        |   Interface    |Proto  |Prio| QL  |Prio| Operational Status |
+-----+------------------------------+----------------+-----+----+-----+----+--------------------+
| F-- |myBITSinput1                  |sync-rj45-1     |BITS |  1 |ST3E |  - |Active              |
| FPT |myPTPinput5                   |PTP_Interface   |PTP  |  - |ST3E |  3 |Active              |
| F-- |mySyncEinput5                 |port 5          |SyncE|  - |ST2  |  1 |Active              |
| F-- |mySyncEinput6                 |port 6          |SyncE|  - |ST2  |  2 |Active              |
| --- |myGPS10MHzInput1              |10mhz-smb-1     |GPS  |  - |PRS  |  - |Active              |
| --- |mySyncEinput8                 |port 8          |SyncE|  - |  -  |  4 |Not Authenticated   |
| --- |myTDMinput5                   |port tdm05      |TDM  |  - |  -  |  - |Loss of Signal      |
+-----+------------------------------+----------------+-----+----+-----+----+--------------------+
```

This example shows sample output for the sync frequency output show command.

```
> sync frequency output show
```

```
+------------------------- FREQUENCY OUTPUT REFERENCES ------------------------+
|                              |     Timing     |       |Oper |                    |
|        Reference Name        |   Interface    |Proto  | QL  | Operational Status |
+------------------------------+----------------+-----+-----+--------------------+
|myGPS10MHzOutput2             |10mhz-smb-1     |GPS  |  -  |Active              |
|myBITSoutput2                 |sync-rj45-1     |BITS |ST3E |Active              |
|myPTPoutput9                  |PTP_Interfac...|PTP  |ST3E |Active              |
|mySyncEoutput1                |port 1          |SyncE|  -  |Loss of Signal      |
+------------------------------+----------------+-----+-----+--------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 5-25
Displaying phase information

You can display:

- phase information
- summary information for phase input references
- summary information for phase output references

Step	Action

To display detailed phase information

1 Display detailed information for phase:

```
sync phase show
```

To display summary information for phase input references

2 Display summary information for input references for phase:

```
sync phase input show
```

To display summary information for phase output references

3 Display summary information for output references for phase:

```
sync phase output show
```

—end—

Copyright© 2022 Ciena® Corporation

Example

This example shows sample output for the sync phase show command.

```
> sync phase show

+-------------------------- PHASE OPERATIONAL STATE --------------------------+
| Parameter              | Value                                              |
+------------------------+----------------------------------------------------+
| Operational Status     | Locked                                             |
| Duration               | 0:02:12:40                                         |
| Operational QL         | ST3E  (last: DNU)                                  |
| Selected Reference     | myPTPinput5                                        |
| Timing Interface       | IP Interface      : PTP_Interface                  |
|                        | Master IP Address : 10.10.10.1                     |
| Protocol               | PTP                                                |
| Ref Switch Count       | 1                                                  |
+------------------------+----------------------------------------------------+

+------------------------------------ PHASE INPUT REFERENCES ------------------------------------+
|Prot |                                |    Timing      |      |Over|Oper |    |                    |
|Group|        Reference Name          |   Interface    |Proto|Prio| QL  |Prio| Operational Status |
+-----+--------------------------------+----------------+-----+----+-----+----+--------------------+
| FPT |myPTPinput5                     |PTP_Interface   |PTP  |  - |ST3E |  3 |Active              |
+-----+--------------------------------+----------------+-----+----+-----+----+--------------------+


+---------------------------- PHASE OUTPUT REFERENCES ----------------------------+
|                                |    Timing      |      |Oper |                    |
|        Reference Name          |   Interface    |Proto| QL  | Operational Status |
+--------------------------------+----------------+-----+-----+--------------------+
|myPTPoutput9                    |PTP_Interfac...|PTP  |ST3E |Active              |
+--------------------------------+----------------+-----+-----+--------------------+
```

This example shows sample output for the sync phase input show command.

```
> sync phase input show

+------------------------------------ PHASE INPUT REFERENCES ------------------------------------+
|Prot |                                |    Timing      |      |Over|Oper |    |                    |
|Group|        Reference Name          |   Interface    |Proto|Prio| QL  |Prio| Operational Status |
+-----+--------------------------------+----------------+-----+----+-----+----+--------------------+
| FPT |myPTPinput5                     |PTP_Interface   |PTP  |  - |ST3E |  3 |Active              |
+-----+--------------------------------+----------------+-----+----+-----+----+--------------------+
```

This example shows sample output for the sync phase output show command.

```
> sync phase output show

+---------------------------- PHASE OUTPUT REFERENCES ----------------------------+
|                                |    Timing      |      |Oper |                    |
|        Reference Name          |   Interface    |Proto| QL  | Operational Status |
+--------------------------------+----------------+-----+-----+--------------------+
|myPTPoutput9                    |PTP_Interfac...|PTP  |ST3E |Active              |
+--------------------------------+----------------+-----+-----+--------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 5-26
Displaying time-of-day information

You can display:

- detailed information for time-of-day
- summary information for input references for time-of-day
- summary information for output references for time-of-day

Step	Action
To display detailed information for time-of-day	
1	Display detailed information for time of day: `sync time-of-day show`
To display summary information for input references for time-of-day	
2	Display summary information for input references for time of day: `sync time-of-day input show`
To display summary information for output references for time-of-day	
3	Display summary information for output references for time of day: `sync time-of-day output show`

—end—

Copyright© 2022 Ciena® Corporation

Example

This example shows sample output for the sync time-of-day show command.

```
> sync time-of-day show

+------------------------ TIME-OF-DAY OPERATIONAL STATE ------------------------+
| Parameter              | Value                                                 |
+------------------------+-------------------------------------------------------+
| Operational Status     | Locked                                                |
| Duration               | 0:02:12:40                                            |
| Operational QL         | ST3E  (last: DNU)                                     |
| Selected Reference     | myPTPinput5                                           |
| Timing Interface       | IP Interface      : PTP_Interface                     |
|                        | Master IP Address : 10.10.10.1                        |
| Protocol               | PTP                                                   |
| Ref Switch Count       | 1                                                     |
+------------------------+-------------------------------------------------------+
| Time-of-day            | Wed Dec 04 23:11:01 UTC 2013                          |
+------------------------+-------------------------------------------------------+

+---------------------------------- TIME-OF-DAY INPUT REFERENCES ------------------------------+
|Prot |                                |    Timing     |      |Over|Oper |     |                |
|Group|        Reference Name          |   Interface   |Proto|Prio| QL  |Prio| Operational Status |
+-----+--------------------------------+---------------+-----+----+-----+----+--------------------+
| FPT |myPTPinput5                     |PTP_Interface  |PTP  |  - |ST3E |  3 |Active              |
+-----+--------------------------------+---------------+-----+----+-----+----+--------------------+

+------------------------- TIME-OF-DAY OUTPUT REFERENCES ------------------------+
|                                |    Timing     |      |Oper |                    |
|        Reference Name          |   Interface   |Proto| QL  | Operational Status |
+--------------------------------+---------------+-----+-----+--------------------+
|myPTPoutput9                    |PTP_Interfac...|PTP  |ST3E |Active              |
+--------------------------------+---------------+-----+-----+--------------------+
```

This example shows sample output for the sync time-of-day input show command.

```
> sync time-of-day input show

+---------------------------------- TIME-OF-DAY INPUT REFERENCES ------------------------------+
|Prot |                                |    Timing     |      |Over|Oper |     |                |
|Group|        Reference Name          |   Interface   |Proto|Prio| QL  |Prio| Operational Status |
+-----+--------------------------------+---------------+-----+----+-----+----+--------------------+
| FPT |myPTPinput5                     |PTP_Interface  |PTP  |  - |ST3E |  3 |Active              |
+-----+--------------------------------+---------------+-----+----+-----+----+--------------------+
```

This example shows sample output for the sync time-of-day output show command.

```
> sync time-of-day output show

+------------------------- TIME-OF-DAY OUTPUT REFERENCES ------------------------+
|                                |    Timing     |      |Oper |                    |
|        Reference Name          |   Interface   |Proto| QL  | Operational Status |
+--------------------------------+---------------+-----+-----+--------------------+
|myPTPoutput9                    |PTP_Interfac...|PTP  |ST3E |Active              |
+--------------------------------+---------------+-----+-----+--------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 5-27
Displaying protection-group information

You can display detailed information for a:

- frequency protection-group
- phase protection-group
- time-of-day protection-group

Step	Action

To display detailed information for a given frequency protection-group

1 Display detailed information for a given frequency protection-group:

```
sync frequency protection-group show group <group>
```

To display detailed information for a given phase protection-group

2 Display detailed information for a given phase protection-group:

```
sync phase protection-group show group <group>
```

To display detailed information for a given time-of-day protection-group

3 Display detailed information for a given time-of-day protection-group:

```
sync time-of-day protection-group show group <group>
```

—end—

Example

This example shows sample output for the sync frequency protection-group show group command, where the group name is myFreqGroup.

```
> sync frequency protection-group show group myFreqGroup

+------------------------ FREQUENCY PROTECTION-GROUP -------------------------+
| Parameter              | Value                                               |
+------------------------+-----------------------------------------------------+
| Group Name             | myFreqGroup                                         |
| Threshold QL           | PROV                                                |
| Forced Reference       | -                                                   |
+------------------------+-----------------------------------------------------+
| Operational Status     | Locked                                              |
| Duration               | 1:01:10:10                                          |
| Operational QL         | ST3E  (last: ST2)                                   |
| Selected Reference     | myBITSinput1                                        |
| Timing Interface       | sync-rj45-1                                         |
| Protocol               | BITS                                                |
| Ref Switch Count       | 2                                                   |
+------------------------+-----------------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

```
+------------------------- FREQUENCY PROTECTION-GROUP INPUT REFERENCES ------------------------+
|Prot |                                |     Timing      |      |Over|Oper|    |                    |
|Group|        Reference Name          |   Interface     |Proto|Prio| QL |Prio| Operational Status |
+-----+--------------------------------+-----------------+-----+----+----+----+--------------------+
| F-- |myBITSinput1                    |sync-rj45-1      |BITS |  1 |ST3E|  - |Active              |
| FPT |myPTPinput5                     |PTP_Interface    |PTP  |  - |ST3E|  3 |Active              |
| F-- |mySyncEinput5                   |port 5           |SyncE|  - |ST2 |  1 |Active              |
| F-- |mySyncEinput6                   |port 6           |SyncE|  - |ST2 |  2 |Active              |
+-----+--------------------------------+-----------------+-----+----+----+----+--------------------+
```

This example shows sample output for the sync phase protection-group show group myPhaseGroup command.

```
> sync phase protection-group show group myPhaseGroup
+--------------------------- PHASE PROTECTION-GROUP ---------------------------+
| Parameter                | Value                                            |
+------------------------+-------------------------------------------------+
| Group Name               | myPhaseGroup                                     |
| Threshold QL             | PROV                                             |
| Forced Reference         | -                                                |
+------------------------+-------------------------------------------------+
| Operational Status       | Locked                                           |
| Duration                 | 0:02:12:40                                       |
| Operational QL           | ST3E (last: DNU)                                 |
| Selected Reference       | myPTPinput5                                      |
| Timing Interface         | IP Interface     : PTP_Interface                 |
|                          | Master IP Address: 10.10.10.1                    |
| Protocol                 | PTP                                              |
| Ref Switch Count         | 1                                                |
+------------------------+-------------------------------------------------+

+-------------------- PHASE PROTECTION-GROUP INPUT REFERENCES ---------------------+
|Prot |                      |     Timing      |      |Over|Oper|    |                  |
|Group|   Reference Name     |   Interface     |Proto|Prio| QL |Prio|Operational Status|
+-----+----------------------+-----------------+-----+----+----+----+------------------+
| FPT |myPTPinput5           |PTP Interface    |PTP  |  - |ST3E|  3 |Active            |
+-----+----------------------+-----------------+-----+----+----+----+------------------+
```

This example shows sample output for the sync phase protection-group show group myToDGroup command.

```
> sync time-of-day protection-group show group MyToDGroup

+------------------------ TIME-OF-DAY PROTECTION-GROUP ------------------------+
| Parameter                | Value                                            |
+------------------------+-------------------------------------------------+
| Group Name               | myToDGroup                                       |
| Threshold QL             | PROV                                             |
| Forced Reference         | -                                                |
+------------------------+-------------------------------------------------+
| Operational Status       | Locked                                           |
| Duration                 | 0:02:12:40                                       |
| Operational QL           | ST3E  (last: DNU)                                |
| Selected Reference       | myPTPinput5                                      |
| Timing Interface         | IP Interface      : PTP_Interface                |
|                          | Master IP Address : 10.10.10.1                   |
| Protocol                 | PTP                                              |
| Ref Switch Count         | 1                                                |
+------------------------+-------------------------------------------------+
| Time-of-day              | Wed Dec 04 23:11:01 UTC 2013                     |
+------------------------+-------------------------------------------------+

+------------------------ TIME-OF-DAY PROTECTION-GROUP INPUT REFERENCES ------------------------+
|Prot |                                |     Timing      |      |Over|Oper|    |                    |
|Group|        Reference Name          |   Interface     |Proto|Prio| QL |Prio| Operational Status |
+-----+--------------------------------+-----------------+-----+----+----+----+--------------------+
| FPT |myPTPinput5                     |PTP_Interface    |PTP  |  - |ST3E|  3 |Active              |
+-----+--------------------------------+-----------------+-----+----+----+----+--------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 5-28
Sample configuration: system timing by means of SyncE and PTP

This sample configuration sets up system timing to:

- recover frequency by means of SyncE (port 5) with PTP (master IP address 2.2.2.1, vlan 200) as backup
- recover phase by means of PTP (master IP address 2.2.2.1, vlan 200)
- distribute frequency by means of BITS (SYNC RJ45 interface) and GPS (10 MHz SMB interface)
- distribute phase by means of GPS (1PPS SMB interface)

The network topology diagram for this configuration is shown in Figure 5-6 on page 5-21.

Step	Action
1	Set the option-type: `sync set option-type option1` ***Note:*** This is a mandatory one-time configuration that must be entered first.

2	Configure the VLAN and IP interface for the PTP session:
	`vlan create vlan 200`
	`vlan add vlan 200 port 2`
	`interface create ip-interface PTP_Interface ip 2.2.2.133/24 vlan 200`

3	Create the PTP input reference object: `sync ptp input create ref PTP_In master-ip-address 2.2.2.1 ip-interface PTP_Interface encap-type udp-over-ip priority 2`
4	Create the SyncE input reference object: `sync synce input create ref SyncE_Port5_In port 5 priority 1`

5	Create the GPS output reference objects:
	`sync gps output create ref SMB_2048_Out gps-interface 10mhz-smb-1 clock-mode frequency frequency-clock 2048khz`
	`sync gps output create ref 1PPS_SMB_Out gps-interface 1pps-smb-1 clock-mode phase`

6	Create the BITS output reference object: `sync bits output create ref E1_BITS_Out bits-interface sync-rj45-1 bits-mode e1 format e1-crc`

Copyright© 2022 Ciena® Corporation

7 Create the frequency protection-group object and add input references with frequency components to it:

```
sync frequency protection-group create group
FreqProtGroup threshold-ql sec
sync frequency protection-group add group FreqProtGroup
input-ref PTP_In,SyncE_Port5_In
```

8 Create the phase protection-group object and add input references with phase components to it:

```
sync phase protection-group create group PhaseProtGroup
threshold-ql sec
sync phase protection-group add group PhaseProtGroup
input-ref PTP_In
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 5-29
Sample configuration: system timing by means of PTP boundary clock

This sample configuration sets up system timing to recover PTP timing on an ordinary clock slave from a grandmaster clock by means of a boundary clock while keeping the grandmaster clock as a backup timing reference.

The network topology diagram for this configuration is shown in "System timing using PTP timing" on page 5-22.

Step	Action

To recover PTP timing on the boundary clock from the grandmaster

1 Configure the VLAN and IP interface for the PTP session:

```
vlan create vlan 200
vlan add vlan 200 port 1, 2
interface create ip-interface ptp200 ip 2.2.2.10/24 vlan
200
```

2 Disable synchronization:

```
sync disable
```

3 Set the synchronization option type to option 2:

```
sync set option-type option 2
```

4 Set the reversion-mode to revertive:

```
sync set reversion-mode revertive
```

5 Set the wait-to-restore value:

```
sync set wait-to-restore 10
```

6 Set the clock type to boundary clock:

```
sync ptp set clock-type bc
```

7 Create the PTP input reference object for the grandmaster clock:

```
sync ptp input create ref GrandMaster encap-type udp-
over-ip ip-interface ptp200 master-ip-address 2.2.2.1
```

8 Create the frequency protection-group object and add input references with frequency components to it:

```
sync frequency protection-group create group
FreqProtGroup threshold-ql tnc
sync frequency protection-group add group FreqProtGroup
input-ref GrandMaster
```


Copyright© 2022 Ciena® Corporation

9 Create the phase protection-group object and add input references with phase components to it:

```
sync phase protection-group create group PhaseProtGroup
threshold-ql tnc
sync phase protection-group add group PhaseProtGroup
input-ref GrandMaster
```

10 Create the time-of-day protection-group object and add input references with time-of-day components to it:

```
sync time-of-day protection-group create group
TodProtGroup theshold-q1 tnc
sync time-of-day protection-group add group TodProtGroup
input-ref GrandMaster
```

11 Create the PTP output reference object:

```
sync ptp output create ref BoundaryClockOutput encap-type
udp-over-ip ip-interface ptp200
```

12 Enable synchronization:

```
sync enable
```

To recover PTP timing on the ordinary clock slave from the boundary clock while keeping the grandmaster as a backup timing reference

13 Configure the VLAN and IP interface for the PTP session:

```
vlan create vlan 200
vlan add vlan 200 port 3
interface create ip-interface ptp200 ip 2.2.2.10/24 vlan
200
```

14 Disable synchronization:

```
sync disable
```

15 Set the synchronization option type to option 2:

```
sync set option-type option 2
```

16 Set the reversion-mode to revertive:

```
sync set reversion-mode revertive
```

17 Set the wait-to-restore value:

```
sync set wait-to-restore 10
```

18 Set the clock type to ordinary clock slave:

```
sync ptp set clock-type oc-slave
```

19 Create PTP input reference object for the boundary clock:

```
sync ptp input create ref BoundaryClock encap-type udp-
over-ip ip-interface ptp200 master-ip-address 2.2.2.10
priority 1
```

20 Create PTP input reference object for the grandmaster clock:

```
sync ptp input create ref GrandMaster encap-type udp-
over-ip ip-interface ptp200 master-ip-address 2.2.2.1
priority 2
```


Copyright© 2022 Ciena® Corporation

21 Create the frequency protection-group object and add input references with frequency components to it:

```
sync frequency protection-group create group
FreqProtGroup threshold-ql tnc
sync frequency protection-group add group FreqProtGroup
input-ref BoundaryClock,GrandMaster
```

22 Create the phase protection-group object and add input references with phase components to it:

```
sync phase protection-group create group PhaseProtGroup
threshold-ql tnc
sync phase protection-group add group PhaseProtGroup
input-ref BoundaryClock,GrandMaster
```

23 Create the time-of-day protection-group object and add input references with time-of-day components to it:

```
sync time-of-day protection-group create group
TodProtGroup theshold-q1 tnc
sync time-of-day protection-group add group TodProtGroup
input-ref BoundaryClock,GrandMaster
```

24 Enable synchronization:

```
sync enable
```

—end—

Copyright© 2022 Ciena® Corporation

Copyright© 2022 Ciena® Corporation

VLAN management

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

These topics are in this section:

- "VLANs and traffic flow" on page 6-1
- "Behavior summary" on page 6-3
- "VLAN/port configuration" on page 6-4
- "VLAN translation" on page 6-5
- "VLAN-based cross-connection" on page 6-6

VLANs and traffic flow

VLANs determine how traffic is forwarded in 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, 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 accepts untagged frames. A port can be configured to admit ONLY tagged frames, admit ONLY untagged frames, or to admit ALL Frames.

Copyright© 2022 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. When the value of the Acceptable Frame Types parameter is set to “VLAN untagged-only”, tagged frames (with single or multiple VLAN tags) received on the port are discarded. In contrast, if the Acceptable Frame Types parameter is set to “all”, all frames are allowed to ingress, regardless of their tag status.

When acceptable frame types are configured on a non-virtual switch UNI, untagged only filters out all three common TPIDs of 0x8100, 0x8898 and 0x9100. For a virtual switch UNI, only the TPID of 0x8100 is recognized as a tagged frame.

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 applies.

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 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. 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, so the VLAN must exist before the corresponding PVID can be assigned. The port does not need to be a member of the VLAN 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

Copyright© 2022 Ciena® Corporation

the ingress port does not belong is 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 port members of VLAN 200, regardless of whether the ingress port belongs to VLAN 200. 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 is 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 has its tag stripped as it egresses. In contrast a frame with a VLAN tag of 300 egresses with its tag intact.

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.

Behavior summary

This table summarizes packet behavior at ingress.

Table 6-1
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.

Copyright© 2022 Ciena® Corporation

Table 6-1
Ingress behavior (continued)

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

This table summarizes packet behavior at egress.

Table 6-2
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 section 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© 2022 Ciena® Corporation

VLAN translation

VLAN tags in a LAN often identify separate logical broadcast domains, for example, a work group. 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 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. Frame VIDs can be either only outer VID or double-tagged VIDs including outer VID and inner VID. The outer VID value is the existing VLAN VID.

Note: MAC learning is performed on the local VLAN.

VLAN translation cannot be configured for:

- a remote management VLAN
- a VLAN that is a member of a multicast service
- a reserved VLAN
- ports that are virtual switch members

Note 1: VS translation shares internal resources with EVPL VS members, MPLS attachment circuits and sub-port instances. This means that the scale of these features may be affected by the number of VLAN translation instances on the switch.

Note 2: The following devices do not support egress VLAN translation for MPLS flows: 5142 and 5160.

Copyright© 2022 Ciena® Corporation

This table shows the number of VLAN translation entries supported per platform.

Table 6-3
Supported VLAN translation entries

Platform	Entries supported
3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928	2048
5142, 5160	8192

This table lists several features related to VLAN functionality and the support provided for VLAN translation.

Table 6-4
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. Classification is based on the raw frame VID, not the translated VID.
IGMP/Multicast services	Supported.
CFM/Y.1731	Not supported.
Link aggregation	Supported.
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.

VLAN-based cross-connection

A layer 2 cross-connect is a point-to-point data path component based on a VLAN feature. There are two types of cross-connections:

- Single VLAN cross-connection — packet forwarding is based on an outer VLAN ID.
- Double VLAN cross-connection — packet forwarding is based on outer and inner VLAN IDs.

Copyright© 2022 Ciena® Corporation

Both single and double VLAN cross-connections have MAC lookup and learning disabled. A packet with a matching cross-connection entry is always treated as a known unicast.

After creating a cross-connection, the existing “vlan set”, “vlan translate” and “vlan queuing” commands still apply to the cross-connection outer-vlan. Bridging VLANs can also be added under cross-connection ports as long as there is no conflict with the outer VLAN IDs.

For hardware data path configurations, a cross-connection entry shares the same resource as MAC entries and OAM endpoints. A maximum of 4096 cross-connections are allowed.

Cross connections support SNMP get and set operations. It is recommended to remove cross connection configurations before deleting its associated VLAN or aggregation port. If a VLAN is used by cross-connection, it does not support other services other than CFM/Y.1731.

TPID configuration

By default, UNI ports support encapsulation on tagged frames with an inner TPID of 0x8100. The network operator can choose to specify a TPID value recognized as a valid TPID value by attachment circuits on the entire platform. This is known as the boxwide inner TPID. Specifying a boxwide inner TPID provides the ability to support different use cases which necessitates the ability to specify the supported inner tag's TPID value.The network operator can choose to set the boxwide inner TPID to

- 0x88a8
- 0x9100

Examples of deployments that make use of boxwide inner TPID configuration are:

- EVPL services where the incoming CTAG with a non-0x8100 TPID value is encapsulated with an STAG that can have a TPID of 0x8100 or 0x88A8 or 0x9100
- EVPL services where the network operator uses a carrier which may expect a specific tag with a TPID value of 0x8100 or 0x88A8 or 0x9100
- an E-ACCESS service where the E-NNI ports support encapsulation of tagged frames having a TPID value of 0x8100 or 0x88A8 or 0x9100

For EVPL services where the incoming CTAG with a non-0x8100 TPID value is encapsulated with an STAG that can have a TPID of 0x8100 or 0x88A8 or 0x9100, 39XX/51XX switches use a virtual switch with an attachment circuit that recognizes the desired CTAG TPID value and a virtual circuit that recognizes and push/pops an STAG with the desired CTAG TPID value, as shown in this figure.

Copyright© 2022 Ciena® Corporation

Figure 6-1
EVPL services: incoming CTAG with a non-0x8100 TPID is encapsulated with STAG of 0x8100 or 0x88A8 or 0x9100.

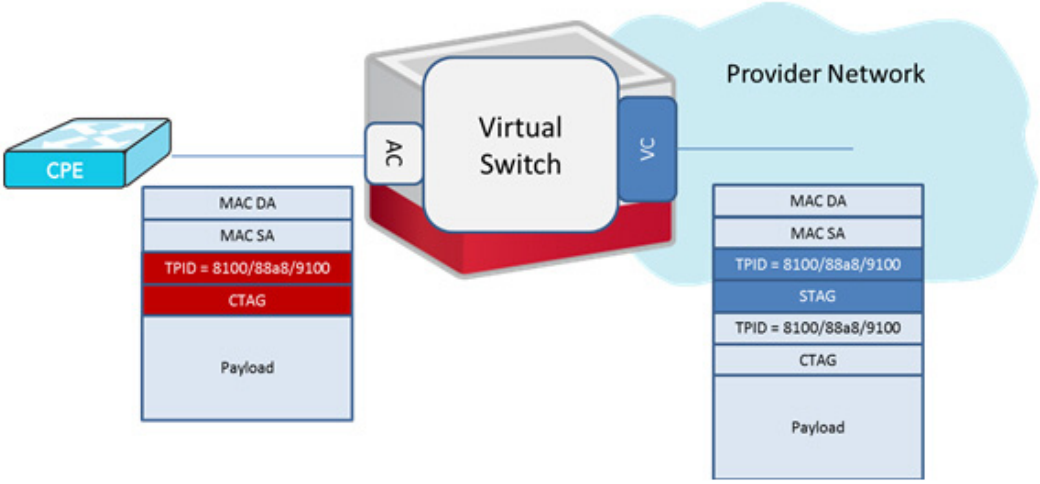

For EVPL services where the network operator uses a carrier which may expect a specific tag with TPID value of 0x8100 or 0x88A8 or 0x9100, 39XX/51XX switches use a VC transform that recognizes and push/pops a PTAG with the desired TPID value, as shown in this figure.

Figure 6-2
EVPL services: network operator uses carrier that expects specific tag

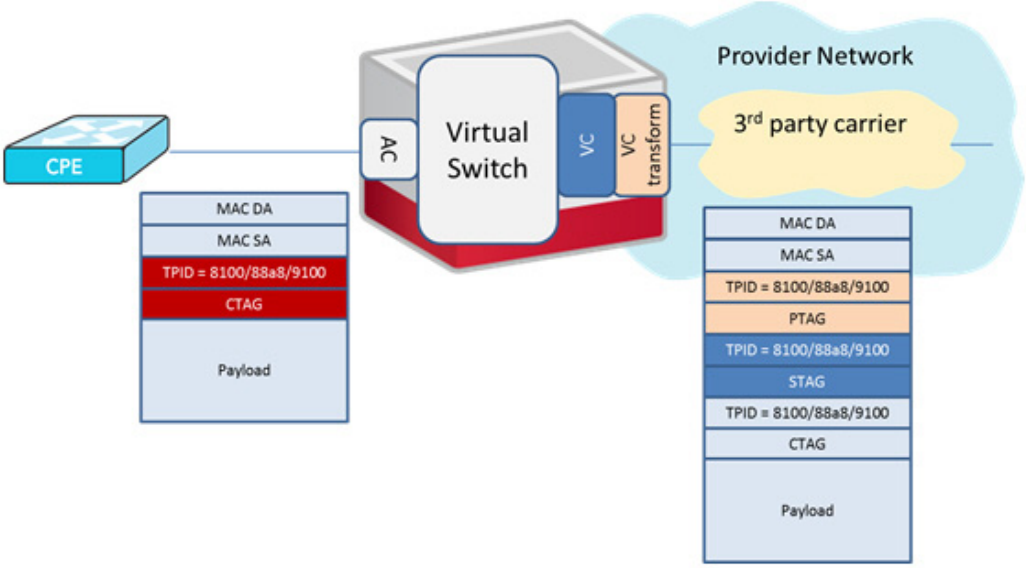

Copyright© 2022 Ciena® Corporation

For an E-ACCESS service where the E-NNI ports support encapsulation of tagged frames which have a TPID value of 0x8100 or 0x88A8 or 0x9100, 39XX/51XX switches stamp the CTAG that has a TPID value recognized by the third-party service provider with an STAG that has the TPID value recognized by the third-party service provider by means of a translate-tag 0 action on an attachment circuit.

The E-ACCESS service: E-NNI ports support encapsulation of tagged frames figure shows an application where the incoming STAG1 is stamped with new STAG2. Each STAG has its own TPID value. If going through a third-party service provider, a PTAG can be added with a VC transform (not shown).

Figure 6-3
E-ACCESS service: E-NNI ports support encapsulation of tagged frames

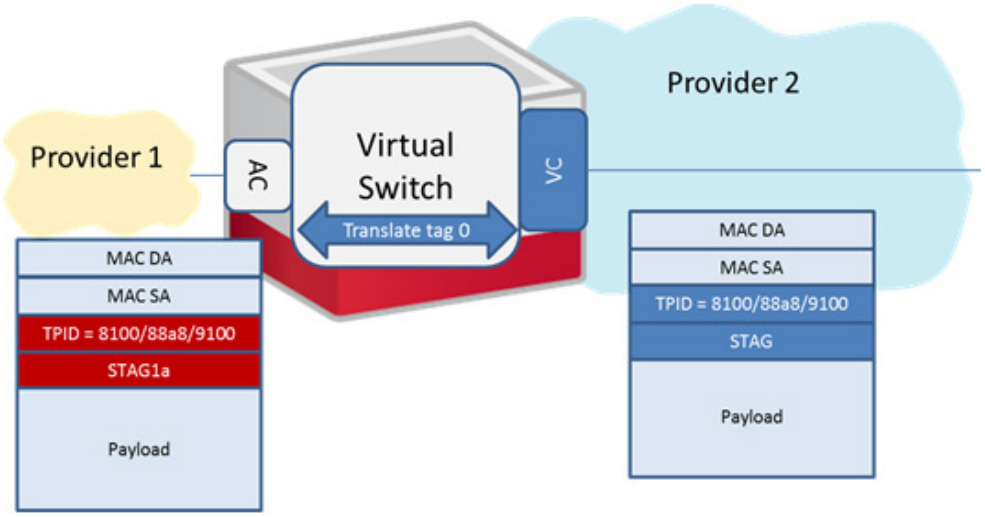

Copyright© 2022 Ciena® Corporation

Procedures

Procedures for VLAN management are:

- “Changing the TPID stamp for a VLAN” on page 6-11
- “Configuring the boxwide inner TPID for a 39XX/51XX switch” on page 6-12
- “Configuring a VLAN/port pair to ingress tagged traffic and egress tagged traffic” on page 6-14
- “Configuring a VLAN/port pair to ingress untagged traffic and egress untagged traffic” on page 6-16
- “Changing tag status” on page 6-18
- “Configuring hybrid traffic” on page 6-20
- “Emulating a tagged Ethernet port” on page 6-22
- “Translating a single NNI VLAN” on page 6-24
- “Translating a dual NNI VLAN” on page 6-26
- “Translating a double-tagged VLAN” on page 6-29
- “Creating a VLAN cross-connection” on page 6-31
- “Deleting a VLAN cross-connection” on page 6-32
- “Displaying VLAN cross-connections” on page 6-33
- “Cross connection deployment example” on page 6-34

Copyright© 2022 Ciena® Corporation

Procedure 6-1
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 port <PortNameList> is the port on the VLAN where the packets egress.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 6-2
Configuring the boxwide inner TPID for a 39XX/51XX switch

By default, UNI ports support encapsulation on tagged frames with an inner TPID of 0x8100. The network operator can choose to set the boxwide inner TPID to

- 0x88a8
- 0x9100

if the boxwide inner TPID is changed to a value other than 0x8100 on a 39XX/51XX switch, the network operator must consider the feature interactions described in this table.

Table 6-5
Feature interactions when TPID is changed from default

Feature	Interaction
Virtual circuit L2 transforms	For an NNI port that is a member of the SVLAN associated with the VC-transform-configured virtual circuit, set the vlan-ethertype attribute to the updated boxwide inner TPID value.
CFM	OAM FPGA-based hardware-assisted CFM and Y.1731 sessions may not function after the boxwide inner TPID is modified. Disable and enable CFM at the global level to restore functionality. For the procedure, refer to *39XX/51XX Switches and Platforms Advanced OAM Configuration*.
Benchmarking	Benchmark generators and reflectors may not work after the boxwide inner TPID is modified. To recover, refer to *39XX/51XX Switches and Platforms Advanced OAM Configuration*.
TDM SFP	TDM SFPs can only use 0x8100 as the boxwide inner TPID value in C-TAG or (S-TAG + C-TAG) frames. These frames may not be identified as single- or double-tagged frames when the boxwide inner TPID is set to any value other than 0x8100.

Copyright© 2022 Ciena® Corporation

Step	Action
1	Change the inner TPID stamp for a VLAN: `vlan set inner-tpid <8100|9100|88a8>` where inner-tpid <8100|9100| 88a8> is the inner TPID. Default value is 8100.
2	Enable the boxwide inner TPID feature: `vlan enable inner-tpid`

—end—

Example

This example configures the boxwide inner TPID to 0x9100 and enables boxwide inner TPID.

```
vlan set inner-tpid 9100
vlan enable inner-tpid
vlan show
+------------ VLAN GLOBAL CONFIGURATION ------------+
| Parameter          | Value                        |
+--------------------+---------------+--------------+
| Inner TPID State   | Enabled                      |
|                    |                              |
| Field              | Admin         | Oper         |
+--------------------+---------------+--------------+
| Inner TPID         | 9100          | 9100         |
+--------------------+---------------+--------------+

+----+--------------------------------+-------------------------+
|VLAN|                                |         111111111122222|
| ID | VLAN Name                 Ports |12345678901234567890123 4|
+----+--------------------------------+-------------------------+
|   1|Default                         | xxxxxxxxxxxxxxxxxx xxxx|
| 127|Mgmt                            | xxxxxxxxxxxxxxxxxx xxxx|
|1000|VLAN#1000                       |xx                      |
|1001|VLAN#1001                       |                        |
+----+--------------------------------+-------------------------+

+----------------------------- CROSS CONNECTION TABLE -------------------------------+
| OVID | IVID |                                Port-A |                            Port-B |
+------+------+-----------------------------------+-----------------------------------+
| No Entry Found                                                                      |
+------+------+-----------------------------------+-----------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 6-3
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 this command: `vlan show [port <PortNameList> | vlan <VlanList>]`

—end—

Copyright© 2022 Ciena® Corporation

Example

This 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© 2022 Ciena® Corporation

Procedure 6-4
Configuring a VLAN/port pair to ingress untagged traffic and egress untagged traffic

This procedure configures a port pair on a VLAN that receives untagged frames and egresses 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 <port> [acceptable-frame-type <all | tagged-only | untagged-only>]` where port <port> — is the port(s) to set. [acceptable-frame-type <all|tagged-only|untagged-only>] — sets the acceptable frame types mode as follows: • all - all frames are allowed to ingress, regardless of their tag status. • tagged-only - only tagged frames are allowed to ingress. • untagged-only - only untagged frames are allowed to ingress. ***Note:*** When acceptable frame types are configured on a non-virtual switch UNI, untagged0only mode filters out all three common TPIDs of 0x8100, 0x8898, and 0x9100. For a virtual switch UNI, only the TPID of 0x8100 is recognized as a tagged frame.
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—**

Copyright© 2022 Ciena® Corporation

Example

This 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© 2022 Ciena® Corporation

Procedure 6-5
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 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>`
3	Set the PVID on the ingress ports to the VLAN used in Step 2. `port set port <PortNameList> [pvid <VlanList>]`
4	To set the ingress fixed .1d priority on the port, enter this command: `port set port port <PortNameList> [fixed-rcos <NUMBER: 0-7>]`

Copyright© 2022 Ciena® Corporation

5 Set the AFT on all ports to All.

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

where

port <port>	is the port(s) to set.
[acceptable-frame-type <all|tagged-only|untagged-only>]	sets the acceptable frame types mode as follows: • all - all frames are allowed to ingress, regardless of their tag status. • tagged-only - only tagged frames are allowed to ingress. • untagged-only - only untagged frames are allowed to ingress. ***Note:*** When acceptable frame types are configured on a non-virtual switch UNI, untagged0only mode filters out all three common TPIDs of 0x8100, 0x8898, and 0x9100. For a virtual switch UNI, only the TPID of 0x8100 is recognized as a tagged frame.

6 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 this example, untagged frames ingressing on port 1 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 egress port 1 untagged.

Assume that port 1 is connected to the subscriber while port 5 connects to the provider network. The frame passes 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, is never seen in the customer's network.

```
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 300
port set port 1 fixed-rcos 5
port set port 1 egress-untag-vlan 300
```


Copyright© 2022 Ciena® Corporation

Procedure 6-6
Configuring hybrid traffic

This procedure configures a port pair to receive either tagged or untagged frames. Tagged frames egress with their original tag intact and untagged frames 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 <port> [acceptable-frame-type <all | tagged-only | untagged-only>]` where port <port> — is the port(s) to set. [acceptable-frame-type <all|tagged-only|untagged-only>] — sets the acceptable frame types mode as follows: • all - all frames are allowed to ingress, regardless of their tag status. • tagged-only - only tagged frames are allowed to ingress. • untagged-only - only untagged frames are allowed to ingress. ***Note:*** When acceptable frame types are configured on a non-virtual switch UNI, untagged0only mode filters out all three common TPIDs of 0x8100, 0x8898, and 0x9100. For a virtual switch UNI, only the TPID of 0x8100 is recognized as a tagged frame.
4	Set the PVID on the ingress ports to VLAN 1 (this is the default setting). `port set port <PortNameList> [pvid <VlanList>]`

Copyright© 2022 Ciena® Corporation

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© 2022 Ciena® Corporation

Procedure 6-7
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 <port> [acceptable-frame-type <all | tagged-only | untagged-only>]` where port <port> — is the port(s) to set. [acceptable-frame-type <all|tagged-only|untagged-only>] — sets the acceptable frame types mode as follows: • all - all frames are allowed to ingress, regardless of their tag status. • tagged-only - only tagged frames are allowed to ingress. • untagged-only - only untagged frames are allowed to ingress. ***Note:*** When acceptable frame types are configured on a non-virtual switch UNI, untagged0only mode filters out all three common TPIDs of 0x8100, 0x8898, and 0x9100. For a virtual switch UNI, only the TPID of 0x8100 is recognized as a tagged frame.

Copyright© 2022 Ciena® Corporation

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 this 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© 2022 Ciena® Corporation

Procedure 6-8
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 <VLAN ID String in outer[:inner] format>} {vlan <NUMBER: 1..4095>}` where port <PortNameList> — is the egress port vid <VLAN ID String in outer[;inner] format> — assign/stamp VIDs in outer[:inner] format. Outer format has to be the same as the VLAN domain. vlan <NUMBER: 1..4095> — assign/stamp to this internal VLAN.
4	Verify the configuration. `vlan translate show port <port>` where port <port> — is the port(s) to show VLAN translation.

—end—

Copyright© 2022 Ciena® Corporation

Example

In the configuration shown in the Single NNI VLAN translation figure, 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.

Figure 6-4
Single NNI VLAN translation

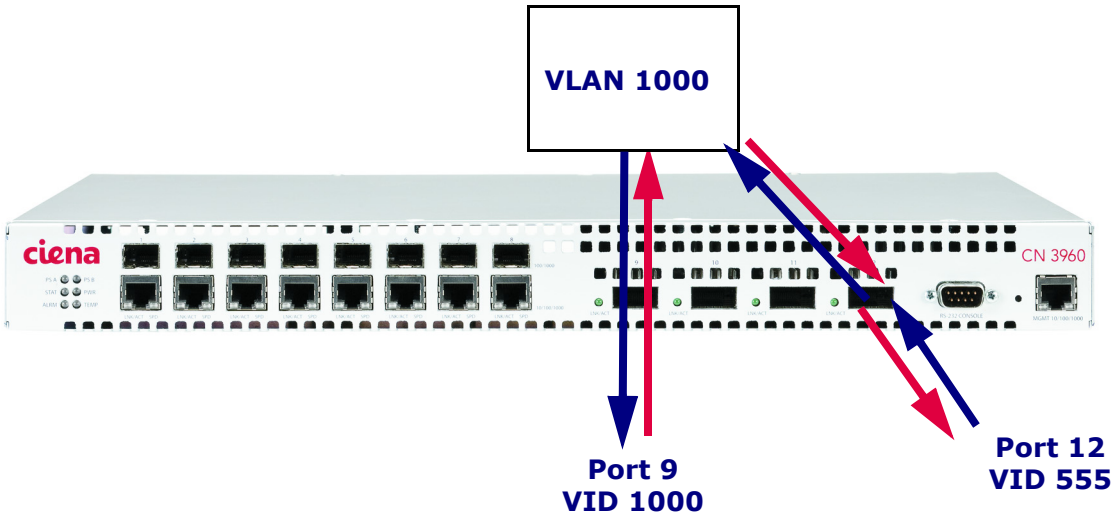

This example creates the configuration.

```
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        | VLAN |                Egress              |
|     |Frame|Frame| To  | To  |      |Internal|Internal|Swap-Frame|Swap-Frame|
|Port|OVID |IVID |OVID|IVID|Domain|  IVID  |  OVID  |IVID      |OVID      |
+----+-----+-----+----+----+------+--------+--------+----------+----------+
|12  |555  |---  |1000|---|1000  |---     | 1000   | ---      |555       |
+----+-----+-----+----+----+------+--------+--------+----------+----------+
```


Copyright© 2022 Ciena® Corporation

Procedure 6-9
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 VID on egress. `vlan translate add port <PortNameList> {vid <VLAN ID String in outer[:inner] format>} {vlan <NUMBER: 1..4095>}` where port <PortNameList> — is the egress port vid <VLAN ID String in outer[;inner] format> — assign/stamp VIDs in outer[:inner] format. Outer format has to be the same as the VLAN domain. vlan <NUMBER: 1..4095> — assign/stamp to this internal VLAN.

Copyright© 2022 Ciena® Corporation

4 Add the translation entry to map the local VLAN to the second VID on egress.

```
vlan translate add port <PortNameList> {<vid <VLAN ID
String in outer[:inner] format>} {vlan <NUMBER: 1..4095>}
```

where

port <PortNameList>	is the egress port
vid <VLAN ID String in outer[:inner] format>	assign/stamp VIDs in outer[:inner] format. Outer format has to be the same as the VLAN domain.
vlan <NUMBER: 1..4095>	assign/stamp to this internal VLAN.

5 Verify the configuration.

```
vlan translate show port <port>
```

where

port <port>	is the port(s) to show VLAN translation.

—end—

Example

In the configuration shown in the Dual NNI VLAN translation figure, 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 VLAN translation entry 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 888.

Copyright© 2022 Ciena® Corporation

Figure 6-5
Dual NNI VLAN translation

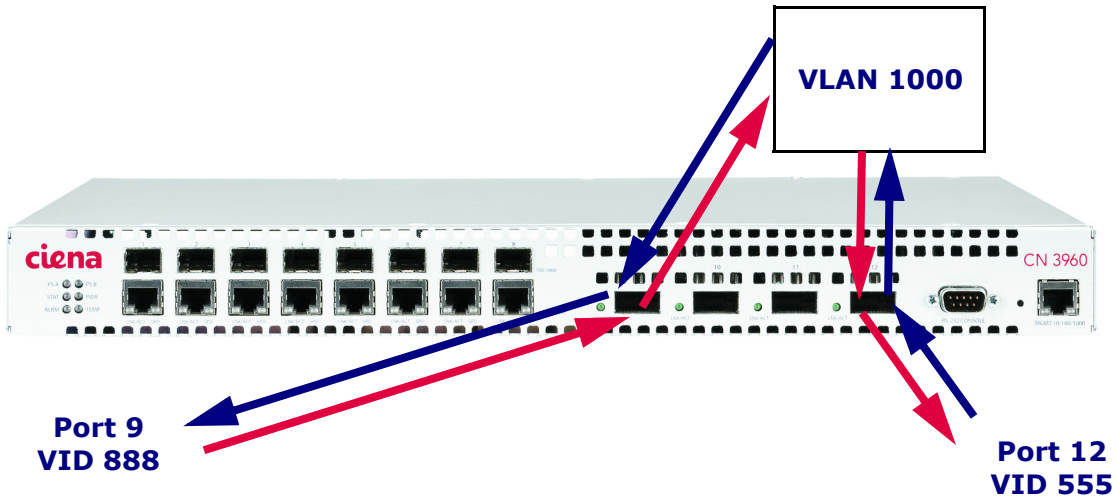

This example creates the configuration.

```
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        | VLAN |                Egress                |
|     |Frame|Frame| To  | To  |      |Internal|Internal|Swap-Frame|Swap-Frame|
|Port |OVID |IVID |OVID|IVID|Domain|  IVID  |  OVID  |  IVID    |  OVID    |
+----+-----+-----+----+----+------+--------+--------+----------+----------+
| 9  |888  |---  |1000|---|1000  |---     | 1000   |---       |888       |
| 12 |555  |---  |1000|---|1000  |---     | 1000   |---       |555       |
+----+-----+-----+----+----+------+--------+--------+----------+----------+
```


Copyright© 2022 Ciena® Corporation

Procedure 6-10
Translating a double-tagged VLAN

Translate a double-tagged 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 <VLAN ID
String in outer:innerformat>} {to-vid <VLAN ID String in
outer:inner format>} {vlan <NUMBER: 1..4095>}
```

where

port <PortNameList>	is the egress port
vid <VLAN ID String in outer:inner format>	assign/stamp VIDs in outer:inner format.
to-vid <VLAN ID String in outer:inner format>	assign/stamp VIDs in outer:inner format. Outer value has to be the same as the VLAN domain.
vlan <NUMBER: 1..4095>	assign/stamp to this internal VLAN.

4 Verify the configuration.

```
vlan translate show port <port>
```

where

port <port>	is the port(s) to show VLAN translation.

—end—

Copyright© 2022 Ciena® Corporation

Example

In the configuration shown in the Double-tagged VLAN translation figure, Port 9 is a member of local VLAN 1000. A double-tagged VLAN translation entry maps frames with VID 123:111 arriving at Port 9 to VID 1000:150. Frames with VID 1000:150 that egress via Port 9 are mapped to VID 123:111.

Figure 6-6
Double-tagged VLAN translation

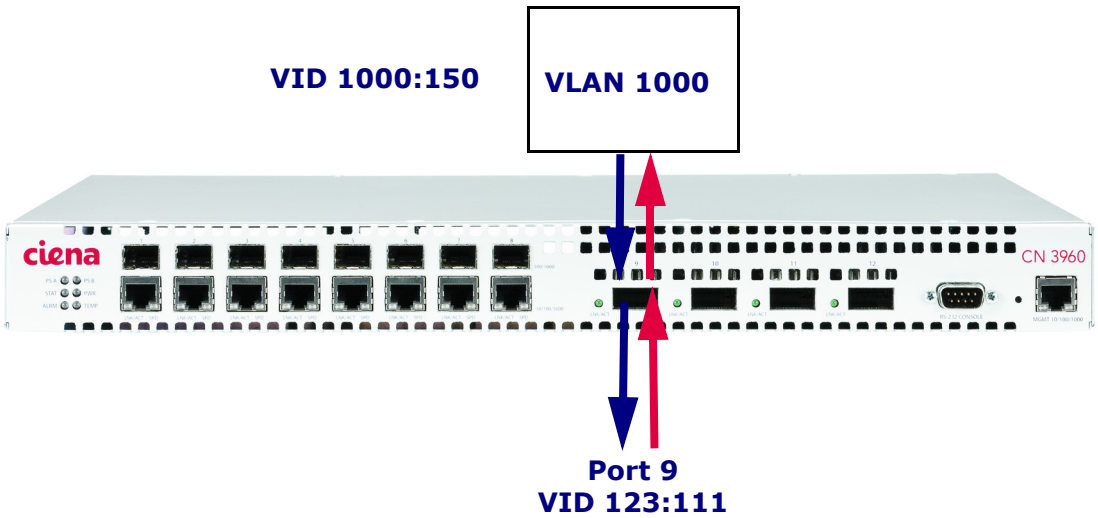

This example creates the configuration.

```
vlan create vlan 1000
vlan add vlan 1000 port 9
vlan translate add port 9 vid 123:111 to-vid 1000:150 vlan 1000
vlan translate show port 9

+---------------------- VLAN TRANSLATE TABLE ----------------------------+
|     |        Ingress        | VLAN |                Egress              |
|     |Frame|Frame| To  | To  |      |Internal|Internal|Swap-Frame|Swap-Frame|
|Port|OVID |IVID |OVID|IVID|Domain|  IVID  |  OVID  |IVID      |OVID      |
+----+-----+-----+----+----+------+--------+--------+----------+----------+
| 9  |123  |111  |1000|150 |1000  |150     | 1000   | 111      |123       |
+----+-----+-----+----+----+------+--------+--------+----------+----------+
```


Copyright© 2022 Ciena® Corporation

Procedure 6-11
Creating a VLAN cross-connection

Create a VLAN cross-connection.

This figure shows the configuration for a simple cross-connection.

Figure 6-7
Cross-connection example

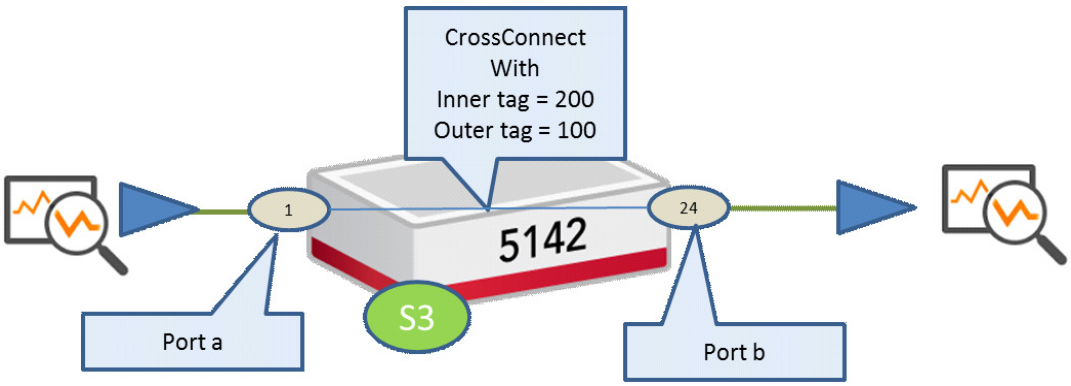

Step	Action
1	Create a VLAN cross-connection:

```
vlan create xc <VLAN ID String in outer[:inner] format>
{port-a <Port>} {port-b <Port>}
```

where

xc <VLAN ID String in outer[:inner] format>	creates a cross-connection with an outer tag (mandatory) and an optional inner tag.
port-a <Port>	is the port on one side.
port-b <Port>	is the port on the other side.

—end—

Example

This example creates a VLAN cross-connection.

```
> vlan create xc 100:200 port-a 1 port-b 24
```


Copyright© 2022 Ciena® Corporation

Procedure 6-12
Deleting a VLAN cross-connection

Delete a VLAN cross-connection.

Step	Action
1	Delete a VLAN cross-connection: `vlan delete xc <VLAN ID String in outer[:inner] format>` where xc <VLAN ID String in outer[:inner] format> — deletes a cross-connection with an outer tag (mandatory) and an optional inner tag.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 6-13
Displaying VLAN cross-connections

Display VLAN cross-connections.

Step	Action
1	Display VLAN cross-connections: `vlan show xc <VLAN ID String in outer[:inner] format>` where xc <VLAN ID String in outer[:inner] format> — displays a cross-connection with an outer tag (mandatory) and an optional inner tag.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 6-14
Cross connection deployment example

The Cross-connection deployment example figure shows an example of a cross-connection deployment. This example shows three single cross-connections and two double cross-connections among different ports. Under the same outer VLAN, either single cross-connection or double cross-connection is supported, but not both.

Figure 6-8
Cross-connection deployment example

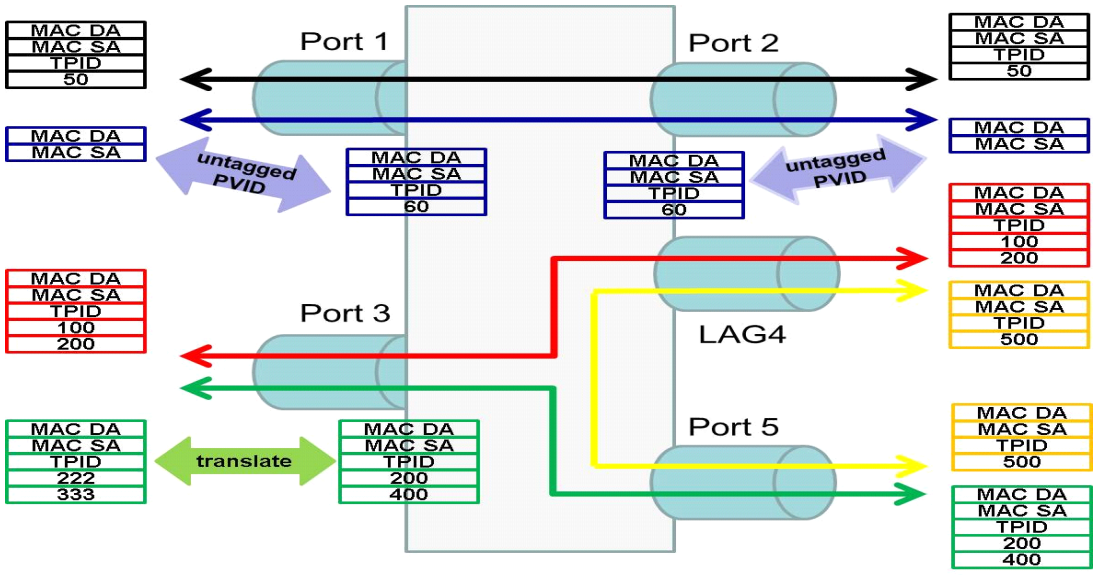

The example consists of these configuration steps:

Copyright© 2022 Ciena® Corporation

Step	Action

Create a single cross-connection

1 Create a single cross-connection matching outer VID 50 between port 1 and 2:

```
vlan create xc 50 port-a 1 port-b 2
vlan show vlan 50
```

```
+--------------------- VLAN 50 INFO ----------------------+
| Parameter             | Value                            |
+-----------------------+----------------------------------+
| VLAN ID               | 50                               |
| Name                  | XC#50                            |
| Features              | XC                               |
| Translation VLAN      |                                  |
| MAC Learning          | Disabled                         |
| Egress TPID           | 8100                             |
| Queue Map*            | Default-RCOS                     |
| PFG State             | Disabled                         |
|     Fwd-Policies      | A -> AB  B -> A                  |
|                                                          |
| *queue map is used only if a queue-group is configured   |
+----------------------------------------------------------+
+--------------------- VLAN Members ----------------------+
| Port       | VTag       | Fwd Group | VS-Sub              |
+------------+------------+-----------+---------------------+
| 1          | 50         | A         | False               |
| 2          | 50         | A         | False               |
+------------+------------+-----------+---------------------+
```

Create a single cross-connection for untagged traffic

2 Create a single cross-connection for untagged traffic:

```
vlan create xc 60 port-a 1 port-b 2
```

Set the ports to accept tagged and untagged traffic

3 Set port 1 and port 2 to accept tagged and untagged traffic. The untagged traffic uses port VID 60.

```
port set port 1,2 acceptable-frame-type all pvid 60
```

Create a LAG port

4 Create a LAG port:

```
aggregation create agg LAG4
aggregation add agg LAG4 port 4,6
aggregation enable
```


Copyright© 2022 Ciena® Corporation

Create a double cross-connection

5 Create a double cross-connection with outer VID 100 and inner VID 200 between port 3 and LAG4:

```
vlan create xc 100:200 port-a 3 port-b LAG4
vlan show xc 100:200
```

```
+---------------------- CROSS CONNECTION TABLE --------------------------+
| OVID | IVID |                    Port-A |                       Port-B |
+------+------+---------------------------+------------------------------+
|  100 |  200 |                         3 |                         LAG4 |
+------+------+---------------------------+------------------------------+
```

Create another double cross-connection

6 Create another double cross-connection with outer VID 200 inner VID 400 between port 3 and port 5:

```
vlan create xc 200:400 port-a 3 port-b 5
```

Create a single cross-connection

7 Create a single cross-connection with outer VID 500 between LAG4 and port 5:

```
vlan create xc 500 port-a LAG4 port-b 5
vlan show xc 500
```

```
+---------------------- CROSS CONNECTION TABLE --------------------------+
| OVID | IVID |                    Port-A |                       Port-B |
+------+------+---------------------------+------------------------------+
| 500  |  --- |                      LAG4 |                            5 |
+------+------+---------------------------+------------------------------+
```

Add a double translation entry for incoming traffic

8 Add a double translation entry for incoming traffic with outer VID 222 and inner VID 333 to outer VID 200 and inner VID 400 on port 3 for the traffic to merge into cross-connection with outer VID 200 and inner VID 400:

```
vlan translate add port 3 vid 222:333 to-vid 200:400 vlan
200
```

```
vlan translate show port 3
```

```
+---------------------- VLAN TRANSLATE TABLE ----------------------------+
|     |        Ingress        | VLAN |                Egress              |
|     |Frame|Frame| To  | To  |      |Internal|Internal|Swap-Frame|Swap-Frame|
|Port|OVID |IVID |OVID|IVID|Domain|  IVID  |  OVID  |IVID      |OVID      |
+----+-----+-----+----+----+------+--------+--------+----------+----------+
| 3  |222  |333  |200 |400 |200   |400     | 200    | 333      |222       |
+----+-----+-----+----+----+------+--------+--------+----------+----------+
```


Copyright© 2022 Ciena® Corporation

Display all VLANs and cross-connections

9 Display all VLANs. A cross-connection table is now shown with the VLAN table:

```
vlan show

+----- VLAN GLOBAL CONFIGURATION -----+
| Parameter        | Value            |
+------------------+------------------+
| Inner TPID       | 8100             |
+------------------+------------------+

+----+--------------------------------+------------+
|VLAN|                                |        111|
| ID | VLAN Name                 Ports |123456789012|
+----+--------------------------------+------------+
|   1|Default                         |xxxxxxxxxxxx|
|  10|VLAN#10                         |xxxxxxxxxxxx|
|  50|XC#50                           |xx          |
|  60|VLAN#60                         |            |
| 100|XC#100                          |  x         |
| 127|Mgmt                            |        xxxx|
| 200|XC#200                          |  x x       |
| 500|XC#500                          |    x       |
|1000|VLAN#1000                       |xxx     x  x|
+----+--------------------------------+------------+

+--------------------- CROSS CONNECTION TABLE -------------------------+
| OVID | IVID |                   Port-A |                       Port-B |
+------+------+--------------------------+------------------------------+
|   50 |  --- |                        1 |                            2 |
|  100 |  200 |                        3 |                         LAG4 |
|  200 |  400 |                        3 |                            5 |
|  500 |  --- |                     LAG4 |                            5 |
+------+------+--------------------------+------------------------------+
```

—end—

Copyright© 2022 Ciena® Corporation

Copyright© 2022 Ciena® Corporation

MEF L2 VPNs

This section describes Layer 2 Virtual Private Networks and how to configure them:

- “Overview” on page 7-1
- “Q-in-Q encapsulation” on page 7-2
- “EPL and EVPL Provider Bridge configuration” on page 7-4
- “Virtual switch emulation of a VLAN on a UNI port” on page 7-10
- “UNI member statistics” on page 7-13
- “EVPL CoS” on page 7-20
- “EVPL CoS” on page 7-20
- “EVPL Bundling” on page 7-21
- “Virtual switch SVLAN transforms” on page 7-25
- “Virtual switch L2 transforms” on page 7-27
- “Virtual circuit transforms” on page 7-34
- “Private forwarding groups” on page 7-45
- “External Network-to-Network Interface Hairpin” on page 7-54
- “MEF services” on page 7-58

Overview

Data moves through a Carrier Ethernet network by means of Point-to-Point and Multipoint-to-Multipoint Ethernet Virtual Connections (EVCs). The types of EVC points are:

- User Network Interface (UNI), which is the physical demarcation between the responsibility of a service provider and the subscriber.

  The customer-side processes of the UNI are called the UNI-C and the network-side processes are called the UNI-N.
- External Network to Network Interface (ENNI), which is the physical demarcation between the responsibility of two service providers. There two types of services providers.

Copyright© 2022 Ciena® Corporation

One type of service provider sells Ethernet services to end user buyers who are connected to the network by means of the UNI. The other type of service provider is the operator who sells Ethernet services to the service provider whose networks are connected by means of the ENNI.

Each service provider is responsible for their respective side of the service demarcation up to the ENNI interconnection point.

The network-side processes on both sides of the ENNI are called ENNI-N.

This figure shows an example of an EVC.

Figure 7-1
EVC example

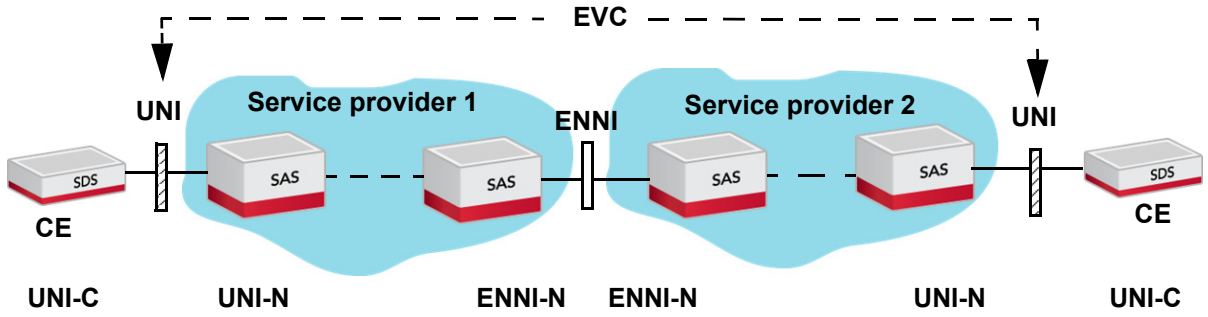

Q-in-Q encapsulation

Support for the MEF service types is configured by means of EPLs and EVPLs that comply with 802.ad Provider Bridges standard (also called Q-in-Q encapsulation and double tagging). Encapsulating the inner 802.1Q Customer VLAN (C-VLAN) tag in the outer 802.1Q Service VLAN (SVLAN) tag enables the service provider to keep subscriber traffic separate, even if the same CVLAN ID is used by more than one subscriber group.

These 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 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 virtual switch of the per port (EPL) member. This can be used in conjunction with the vs-ingress-filter.

Copyright© 2022 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.

Three UNI port options impact how EPL and EVPL treats tagged and untagged traffic:

- vs-ingress-filter
- untagged-data-vs
- untagged-ctrl-vs

This table summarizes ingress filtering.

Table 7-1
Ingress filtering

UNI port attribute	Untagged frames	Tagged frames
untagged-data-vs <vs>	A virtual switch EPL member for the UNI port does not accept untagged data frames unless the untagged-data-vs attribute on the UNI port is set to that virtual switch. The virtual switch specified by <vs> must contain the EPL member of that port.	Not applicable

Copyright© 2022 Ciena® Corporation

Table 7-1
Ingress filtering

UNI port attribute	Untagged frames	Tagged frames
untagged-ctrl-vs <vs>	A VS EPL member for that UNI port does not accept untagged control frames unless the untagged-ctrl-vs on the UNI port is set to that VS. The virtual switch specified by <vs> must contain the EPL member of that port.	Not applicable
vs-ingress-filter on	Not applicable	The EVPL member for the outer VID of the frame must be defined for the UNI port in order for that frame to be processed by that EVPL member. If the EVPL member for the outer VID of the frame is not defined for the UNI port, the EPL member for the UNI port does not process that frame.
vs-ingress-filter off	Not applicable	The EVPL member for the outer VID of the frame on the UNI port processes that frame. If the EVPL member for the outer VID of the frame is not defined for the UNI port, the EPL member for that UNI port, if defined, processes that frame.

Note: If the vs-ingress-filter of the UNI port is set to on, the UNI EPL of a virtual switch does not accept tagged data frames for which the outer VID in the frame does not match an EVPL member on the UNI port.

EPL and EVPL Provider Bridge configuration

EPL and EVPL Provider Bridge configuration is implemented by attaching a virtual circuit with its associated Provider VLAN, to a virtual switch, which defines the subscriber ports (PP) or subscriber ports and VLANs (PPV). From the subscriber's perspective, the subscriber connection to the service provider network is a direct connection using a private LAN between geographically-separated sites.

The general steps for configuring an EPL or EVPL are:

1 Create a service provider VLAN.
2 Add the NNI port to the service provider VLAN.
3 Create an Ethernet virtual circuit using the service provider VLAN.

Copyright© 2022 Ciena® Corporation

4 Create the virtual switch and associate it with the virtual circuit. This is an optional step that is only required if the Ethernet virtual circuit is not specified when creating the Ethernet virtual switch.

5 Add the UNI members to the Ethernet virtual switch.

Figure 7-2
EPL and EVPL configuration overview

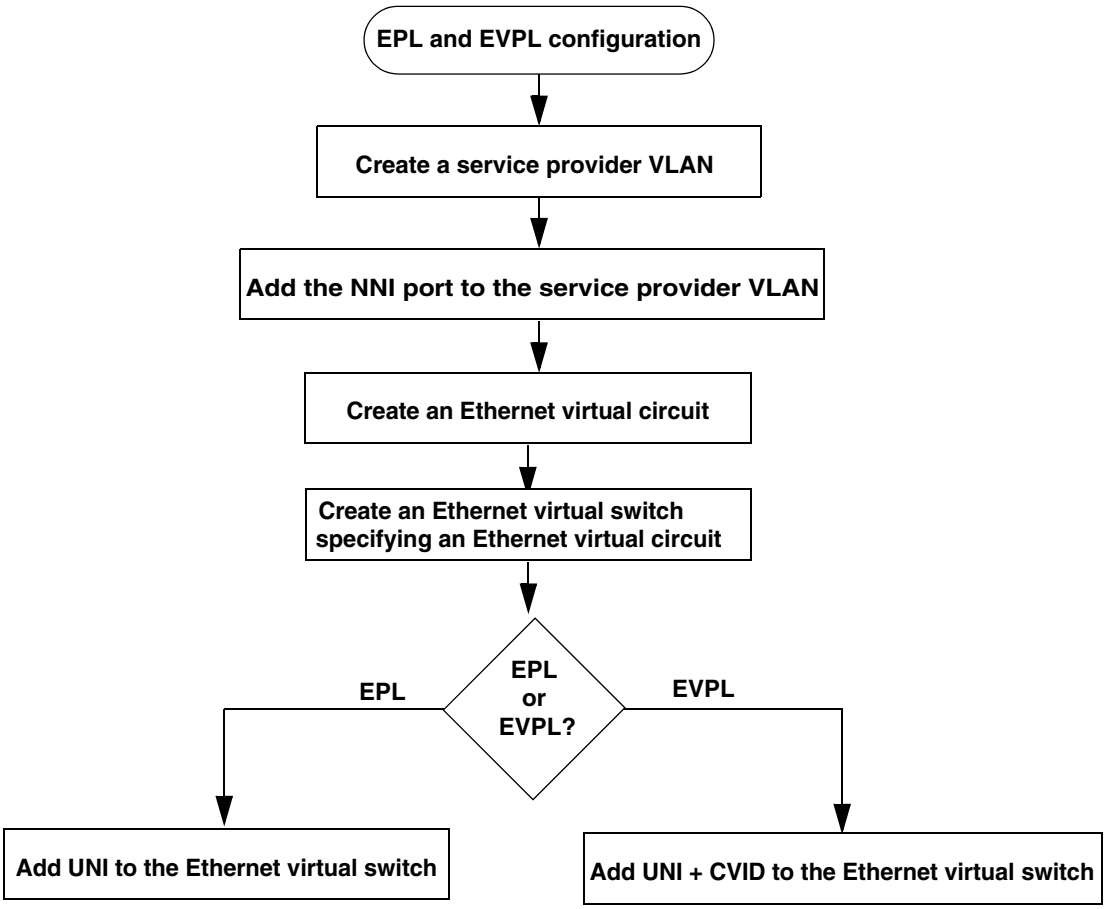

Provider VLAN

In the fully configured EPL or EVPL, the provider VLAN is the Service VLAN (SVLAN) outer 802.1D tag and has an NNI (egress) 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.

This figure shows a provider VLAN overview using Q-in-Q.

Copyright© 2022 Ciena® Corporation

Figure 7-3
Provider VLAN overview using Q-in-Q

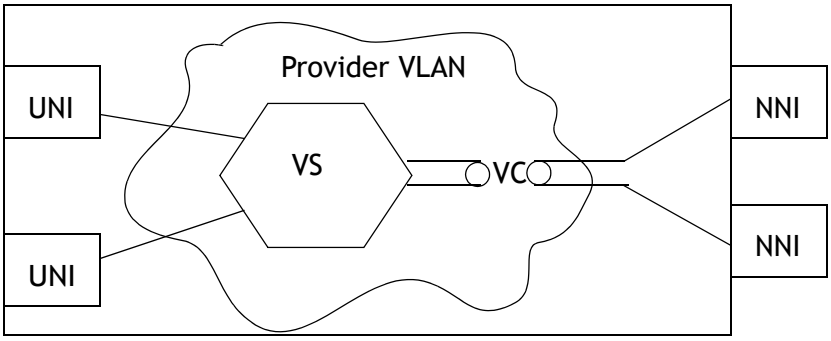

Reserved VLAN

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 created with these exceptions:

- The virtual switch is VPLS or VPWS.
- The virtual switch is created as an Ethernet Q-in-Q, that is its Ethernet virtual circuit is assigned when the virtual switch is created.

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 virtual switch is created or you can specify a reserved VLAN when you create the virtual switch.

Subscriber/UNI traffic can also be switched without connecting to the provider/transport network.

If a reserved VLAN is specified while an Ethernet virtual switch is created, it uses the reserved VLAN as its L2 switching domain while an Ethernet (Q-in-Q) virtual circuit is not assigned to the virtual switch. If a reserved VLAN is not specified while an Ethernet virtual switch is created, the Ethernet virtual circuit must be specified in the same command.

When a VLAN is added to the pool of reserved VLANs, it uses one of the available 4092 VLANs. A reserved VLAN can only be used for Q-in-Q VS. If an Ethernet VS has an attached Ethernet VC and a specified reserved VLAN, the reserved VLAN is not utilized.

A reserved VLAN (r-VLAN) can be removed from an Ethernet virtual switch only if it is currently assigned an Ethernet virtual circuit.

An r-vlan can be freed from a virtual switch if specified in the VS create command using the unset command.

Copyright© 2022 Ciena® Corporation

During a software upgrade from a load where Ethernet virtual switches are always assigned a reserved VLAN (prior to SAOS 6.13), all Ethernet virtual switches configured include a configured reserved VLAN. During a software downgrade to a load where Ethernet virtual switches are always assigned a reserved VLAN (prior to SAOS 6.13), any virtual switches that are not configured with a reserved VLAN and no reserved VLANs are available from the pool of reserved VLANS, the virtual switch and its members are removed from the configuration file.

This figure shows a reserved VLAN.

Figure 7-4
Reserved VLAN

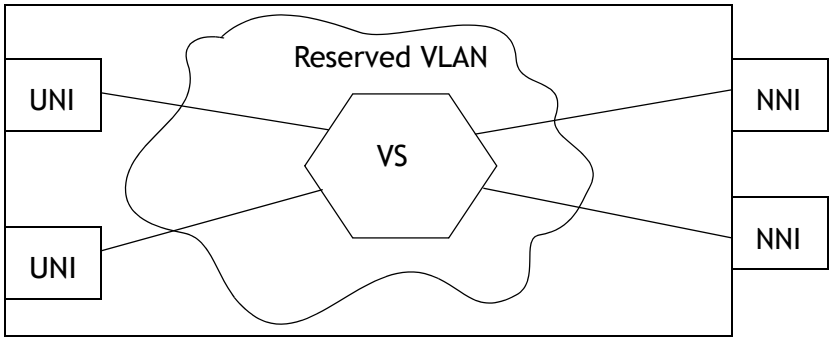

Virtual switch

A virtual switch is a logical entity that can co-exist with other virtual switches in a physical device. Virtual switches allow the physical device to provide a variety of methods for classification, segregation, and flexibility in frame processing. Virtual switches segment the system into separate logical forwarding domains or flood domains to enable VLAN identifiers (VIDs) to be reused across multiple virtual switches. For example, a VID of 101 instanced on two or more virtual switches with unique port connectivity to each virtual switch allows multiple customers to use VID 101 for their own network connectivity, yet be completely isolated from each other. A virtual switch provides the Layer 2 forwarding domain between its UNI members and the NNI. It provides the ability to bring together different networks, while maintaining their security and isolation.

Virtual switches are the switching construct used for Q-in-Q (as described in this section), as well as, for MPLS and PWE.

> ***Note:*** When a port is added to a virtual switch, that port is automatically removed from its external VLANs that port was a member of. This configuration change is added to the configuration file.

Copyright© 2022 Ciena® Corporation

The number of virtual switches and members supported for Q-in-Q depends on the platform capabilities as shown in this table.

Table 7-2
Q-in-Q virtual switch capabilities

Platform	Virtual switches	Virtual switch members for each port	Virtual switch members for each switch
3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928	1024	2048	2048
3942	4092	4095	8192
5142, 5160	4092	4095	16384

Note 1: The number of virtual switches and virtual switch members for each port and switch supported for Q-in-Q are reduced by the number of virtual switches and virtual switch members for each port and switch configured for MPLS.

Note 2: Virtual circuits and virtual switches have a 1:1 relationship. A virtual circuit can only be associated with one virtual switch. A virtual switch can only be associated with one virtual circuit.

Virtual circuits

A virtual circuit defines a secure logical connection between one or more customer endpoints. It is a service trunk, where transport services start and end. Virtual circuits use the SVLAN (or Provider VLAN) to tag traffic from the UNI side (encapsulate) and strip the SVLAN (or Provider VLAN) from traffic coming from the NNI side (decapsulate) 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.

The number of virtual circuits supported for Q-in-Q depends on the platform capabilities as shown in this table.

Table 7-3
Q-in-Q virtual circuits supported by platform

Platform	Virtual circuits
3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928	1024
3942, 5142, 5160	4092

Copyright© 2022 Ciena® Corporation

SAOS allows Ethernet virtual circuits to be statically created. An Ethernet-virtual circuit uses an additional 802.1Q VLAN tag and includes:

- Service Etype (default 0x8100)
- Service VLAN ID (value 1-4094)
- Service VLAN Priority Code Point (PCP) (also, called SVLAN 802.1D priority) (value 0-7)

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 is flooded to:

- 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 CVID of the frame matches the CVID configured against the EVPL member

Q-in-Q Ethertype

By default the outer tag TPID of frames egressing a VLAN member 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.

This setting supports 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 this table.

Table 7-4
Ethertype setting by 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 the Ethertype policy by port table.
88a8	This indicates to use the 802.1ad Ethertype on frames using the user configured policy described in the Ethertype policy by port table.

Copyright© 2022 Ciena® Corporation

The `virtual-circuit ethernet set port` command includes a setting to specify the policy or rules for applying that configured Ethertype to each port as described in this table.

Table 7-5
Ethertype policy by port

Name	Description
all	When this policy is set, all frames going out the port are stamped with the configured Ethertype (see the Ethertype setting by port table).
vlan-tpid	When this policy is set, frames going out a port are stamped with the TPID value of the egress-tpid value assigned to the provider VLAN.

Virtual switch emulation of a VLAN on a UNI port

Using Ethernet virtual switching and a virtual switch transform mode, it is possible to emulate VLAN switching for certain frames on the same UNI port that also has frames processed by means of a per-port (PP) Q-in-Q or per-port-per-VLAN (PPV) Q-in-Q configuration.

Copyright© 2022 Ciena® Corporation

This table lists the tasks required to support this configuration:

Table 7-6
Virtual switch emulation of a VLAN on a UNI port

Procedure	Description
Setting the virtual switch L2 transform	The vs-l2-transform mode of the UNI port must be set to i-stamp;push, e-match-pop:stamp. This allows virtual switch members on the UNI port to use the translate-tag capability including the translate tag zero function. This replaces the CVID of the frame with an S-VID when ingressing the virtual switch member. the opposite action occurs when egressing the virtual switch member.
Adding a Q-in-Q EVPL service on the UNI port	Create a VLAN to represent the SVLAN and add NNI ports as members.
	Create an Ethernet virtual circuit using the SVLAN.
	Create an Ethernet virtual switch using the Ethernet virtual circuit.
	Create an EVPL virtual switch member on the UNI port and, • If CVID translation is not required, specify the CVID as the translate-tag value or do not specify the translate-tag. • If CVID translation is required, specify the incoming CVID to be translated as the translate-tag value. See "Virtual switch L2 transforms" on page 7-27 for more information on translate-tag usage. ***Note:*** The configuration steps required to add a QiQ EPL service on the UNI port are the same regardless of the vs-l2 transform mode on the port.

Copyright© 2022 Ciena® Corporation

Table 7-6
Virtual switch emulation of a VLAN on a UNI port

Procedure	Description
Emulating VLAN switching on the UNI	Create a VLAN representing the VID to the switched traffic and add NNI ports as members.
	Create an Ethernet virtual circuit using the VLAN created.
	Create an Ethernet virtual switch using the Ethernet virtual circuit.
	Create an EVPL virtual switch member on the UNI port, with the vlan/CVID equal to the VID of the traffic you are switching, and a translate-tag of zero.
	Note: Specifying an EVPL virtual switch member with a translate-tag of zero stamps the incoming VID of the frame (CVID) with the VID of the virtual circuit (S-VID), and on the egress stamp, the incoming VID of the virtual circuit (S-VID) with the VID of the EVPL (CVID). Since the EVPL VID (CVID) and VLAN of the VC (S-VID) are the same, there is no change to the VID of the frame and it appears to be VLAN switched based on the VID.

An example of using a virtual switch to emulate a VLAN is to have remote management VLAN traffic switched through a UNI port of the device, where the UNI port also serves customer traffic for a Q-in-Q network. This allows downstream devices attached to Q-in-Q network access nodes to be managed using the remote management VLAN.

The Aggregation network example figure shows three downstream devices connected to an aggregate device. Each downstream device is connected via one port that transports both customer data traffic and device management traffic.

Copyright© 2022 Ciena® Corporation

Figure 7-5
Aggregation network example

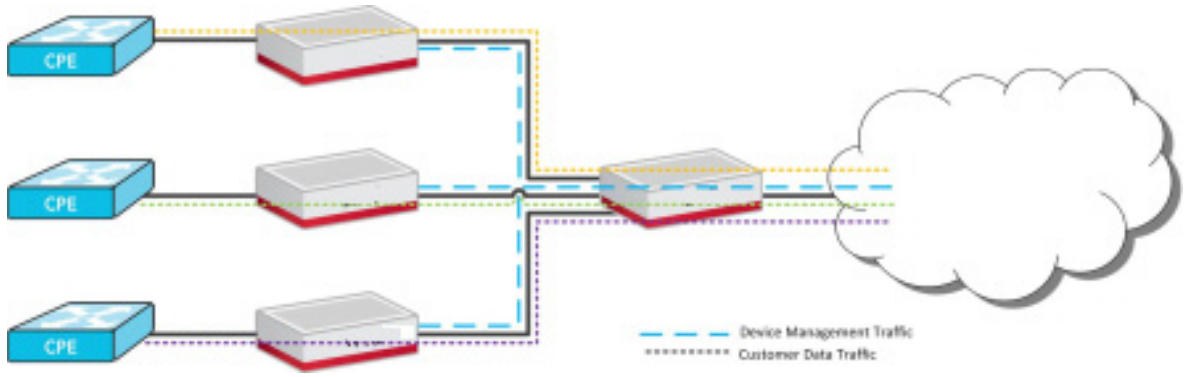

Translate-tag zero and double tagged UNI frames

The user can configure a virtual switch with a translate-tag zero service to emulate a VLAN. In general, this service in intended to switch a single tagged UNI port frame to a single tagged NNI port frame. It may be tempting to use the same service definition to switch a double tagged UNI port frame to a double tagged NNI port frame.

Double tagged traffic from UNI to NNI direction on a translate-tag zero member cannot have its inner PCP preserved. Instead, it is copied from the outer VLAN tag because of a SAOS limitation.

The VLAN emulation service can be used to switch a double tagged frame, but the following limitations apply when virtual switch L2 enhanced transform resources are allocated:

> ***Note:*** Virtual switch L2 enhanced transform resources are allocated to the applicable devices when configuring Ethernet virtual private lines (EVPLs) with non-zero translate-tag attributes.

- When these resources are allocated, egress UNI traffic on a translate-tag zero service has outer PCP and DEI bit values based on the inner VLAN tag.
- When these resources are not allocated, the PCP and DEI bit values of egress UNI traffic on a translate-tag zero service are not affected and are based on the outer VLAN tag that ingressed the NNI port.

UNI member statistics

A port can be added in a virtual switch as a UNI member with a specific VLAN (EVPL Type) and also without specifying any VLAN (EPL Type). Statistics collection for UNI members can be configured on or off depending on resource limits.

This figure shows ingress meter function represents the traffic-profiles configured for each UNI member of the virtual switch.

Copyright© 2022 Ciena® Corporation

Figure 7-6
UNI member statistics example

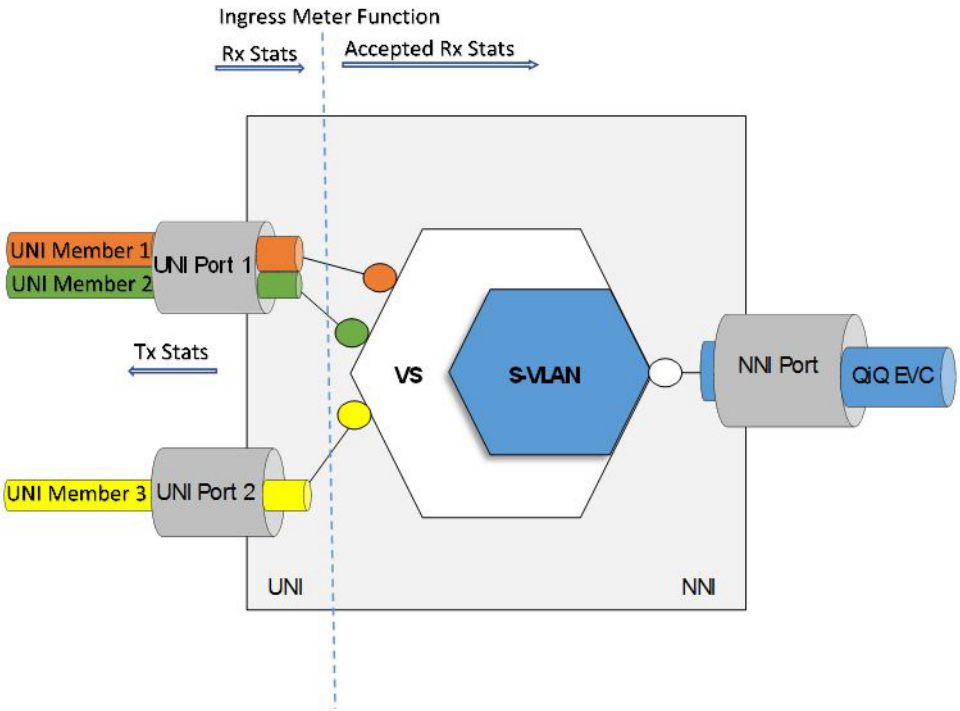

These types of statistics are provided for UNI members of a QinQ virtual switch:

- Packet and byte counts at ingress of UNI member (Rx statistics) for the traffic flow from UNI to NNI direction.
- Packet and byte counts after ingress bandwidth profile (accepted Rx statistics) for the traffic flow from UNI to NNI direction.
- Packet and byte counts egressing a UNI member (Tx statistics) for the traffic flow from NNI to UNI direction.

Traffic profiles

UNI member statistics depends on traffic profiling for Rx statistics and accepted Rx statistics. It is recommended that you configure a traffic profile for each UNI member of a virtual switch. Traffic profiles are not automatically created for a UNI member with statistics collection turned on.

If the statistics are turned on and the traffic profile is already provided for the UNI member, you can set a different traffic profile on the same UNI member without turning off the statistics collection. Statistics can be turned on without providing a traffic profile.

Copyright© 2022 Ciena® Corporation

Traffic profiles being used for statistics collection can be modified and cleared by the user.

This table describes traffic profile behavior.

Table 7-7
Traffic profile behavior

Traffic profile is	Behavior
Provided	The device does not perform any validation to verify the traffic profile provided for the UNI member statistics. The device uses the traffic profile's statistics to calculate Rx statistics and accepted Rx statistics of the UNI member.
Unset (removed)	The device provides Tx statistics. It does not update Rx statistics, and it does not update accepted Rx statistics.
Not provided	The device provides Tx statistics. There is no update in Rx statistics. There is no update in accepted Rx statistics.
Deleted	You receive a warning message. The traffic profile is deleted from the device. There is no update in Rx statistics and accepted Rx statistics.
Globally disabled	You receive a warning message. There is no update in Rx statistics and accepted Rx statistics.
Disabled for a port	You receive a warning message. There is no update in Rx statistics and accepted Rx statistics.

Copyright© 2022 Ciena® Corporation

Statistics and statistics resources

This table describes the types of statistics collected.

Table 7-8
Statistics collected

Statistics	Description
Rx bytes	Collected for the UNI member with statistics collection turned on. Rx bytes is the number of frame bytes received by the UNI member. The device uses the sum of accepted and dropped bytes count of the traffic profile configured for the UNI member to provide Rx bytes count.
Rx packets	Collected for the UNI member with statistics collection turned on. Rx packets is the number of frames received by the UNI member. The device uses the accepted and dropped packets count if the traffic profile configured for the UNI member to provide Rx packets count.
Accepted Rx bytes	Collected for the UNI member with statistics collection turned on. Accepted Rx bytes is the number of frame bytes accepted by the traffic profile configured for the UNI member. The device uses the accepted bytes count of the traffic profile configured for the UNI member to provide accepted Rx bytes count.

Copyright© 2022 Ciena® Corporation

Table 7-8
Statistics collected

Statistics	Description
Accepted Rx packets	Collected for the UNI member with statistics collection turned on. Accepted Rx packets is the number of frames received by the UNI member. The device uses accepted and dropped packets count of the traffic profile configured for the UNI member to provide Rx packet count.
Tx bytes	Collected for the UNI member with statistics collection turned on. Tx bytes is the number of frames transmitted by the UNI member.
Tx packets	Collected for the UNI member with statistics collection turned on. Tx packets is the number of frames transmitted by the UNI member.

A maximum of 256 UNI members supports statistics on the 39xx devices. whereas on 51xx devices, the scale can reach up to 512.

Each UNI member which is collecting statistics uses statistics resources for Tx statistics collection as follows:

- For each EPL or EVPL member with statistics collection turned on, one statistics resource is used.
- Two additional resources are used for an EVPL member where untagged statistics are turned on along with statistics collection.
  - One resource is used to count untagged frames transmitted out of a given EVPL member.
  - One resource is used to count priority tagged frames transmitted out of a given EVPL member.

When a port is added as an EVPL member of the virtual switch with statistics collection turned on, then the statistics for untagged traffic transmitted out of that port are added with the Tx statistics of the same UNI member if these conditions are met:

- Untagged statistics collection is enabled on the same UNI member.
- The same port is not added in any virtual switch as an EVPL member with statistics collection turned on along with untagged statistics enabled.

Copyright© 2022 Ciena® Corporation

- The same port is not added in any virtual switch as an EPL member with statistics collection turned on, along with untagged statistics enabled.

When a port is added as an EPL member of the virtual switch with statistics collection turned on, then the statistics for untagged traffic transmitted out of that port are added with the Tx statistics of the same UNI member if these conditions are met:

- Untagged statistics collection is enabled on the same UNI member.
- The same port is not added in any virtual switch as an EVPL member with statistics collection turned on along with untagged statistics enabled.

Feature interaction

UNI member statistics depend on the traffic profiling feature for received statistics and accepted received statistics. It is recommended that the traffic profile for each UNI member of a VS be configured.

Standard profiles which are created using the traffic-profiling “profile-mode” as port can only be used for UNI member’s statistics collection. These procedures are when a traffic profile is configured for an

- EVPL member of a virtual switch
- EPL member of a virtual switch
- EVPL member of a virtual switch with untagged statistics
- EPL member of a virtual switch with untagged statistics

Caveats

UNI member statistics introduce these caveats:

- A frame which is counted as an accepted frame in accepted Rx statistics may get dropped because of L2CFT, Broadcast Containment, Access Control List or Egress queue congestion.
- Traffic profiles must be configured for each UNI member to get the correct offered Rx and accepted Rx statistics.
- Statistics collection is not supported for UNI members of MPLS-based virtual switches.
- Statistics collection is not supported for TDM port-based UNI members.
- Statistics collection is not supported for sub-port-based UNI members.
- Priority tagged frames ingressing on a port can only be qualified by EPL-based traffic profiles whereas priority tagged frames egressing out of the same port are counted as part of the virtual member where “include-untagged-statistics” flag is set to “on”.

Copyright© 2022 Ciena® Corporation

- A port with both EPL and EVPL memberships of a single virtual switch or multiple virtual switches and the EPL member's statistics collection is enabled, then the Tx statistics of the EPL member also include statistics of all EVPL members where statistics collection is not enabled. There are two possible workarounds for this limitation:
  - The EPL member's statistics collection is only enabled when statistics collection is enabled for all EVPL members of the same port.
  - Only enable statistics collection for specific EVPL members. Do not enable statistics collection for the EPL member.

Use case

There can be a use case where statistics are required per EVC basis instead of a per UNI member basis. In this case, the sum of a UNI member's statistics can be used to calculate the EVC statistics.

The example in this table can be used to get throughput statistics using the virtual switch member's statistics.

Table 7-9
Use case example

Example 1	Statistics
At time T1	• Get Rx Bytes (Rx1) • Accepted Rx Bytes (ARx1) • Tx Bytes (Tx1) statistics • Time stamp (TS1) in hundredths of a second from the device.
At time T2	• Get Rx Bytes (Rx2) • Accepted Rx bytes (ARx2) • Tx Bytes (Tx2) Statistics • Time stamp (TS2) in hundredths of a second from the device. • Rx Throughput (bps)=((Rx2-Rx1) * 8)/((TS2-TS1)/100) • Accepted Throughput (bps)=((ARx2-ARx1)*8)/((TS2-TS1)/100) • Tx Throughput (bps)=((Tx2-Tx1) * 8)/((TS2-TS1)/100)

Note: Time stamp is provided by device for better accuracy when calculating throughput.

Copyright© 2022 Ciena® Corporation

See these procedures:

- “Configuring UNI member statistics” on page 7-71
- “Clearing UNI member statistics” on page 7-75
- “Displaying UNI member information” on page 7-76

EVPL CoS

Class of Service is configured by means of:

- “Virtual switch CoS policies”
- “Virtual switch member CoS policy override”

Virtual switch CoS policies

Each virtual switch 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 virtual switch is created the default encap-fixed-dot1dpri value is 2 and the subscriber-dot1dpri-policy is leave.

Virtual switch member CoS policy override

By default, virtual switch members inherit the CoS policy from the virtual switch. You can configure the encap-cos-policy of the virtual switch member to one of three settings:

- fixed — The SVLAN 802.1D value is taken from the encap-fixed-dot1dpri parameter of the virtual switch member.
- port-inherit — The SVLAN 802.1D value is based on the resolved Cos policy of the port, which can be one of (see “Class of Services (CoS) policies and mapping” on page 8-1 to understand the restrictions for other EVPL members when one or more EVPL members override the VS CoS policy with the port-inherit setting):
  - 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 this table.

Table 7-10
Default .1D mapped CoS mapping table

Customer Frame 802.1D	SVLAN 802.1D
0	0
1	1
2	2

Copyright© 2022 Ciena® Corporation

Table 7-10
Default .1D mapped CoS mapping table

Customer Frame 802.1D	SVLAN 802.1D
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 this table.

Table 7-11
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 virtual switch encapsulation CoS policy configuration. This setting is the default.

EVPL Bundling

SAOS supports EVPL bundling. EVPL bundling allows multiple customer VLAN IDs (CVIDs) to be mapped to one or more virtual switches. Note that EVPL bundling only functions where traffic is being encapsulated.

As an example, the Multiple CVIDs to multiple virtual switches figure shows multiple CVIDs 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 CVIDs to be mapped to two separate virtual switches.

Copyright© 2022 Ciena® Corporation

Figure 7-7
Multiple CVIDs to multiple virtual switches

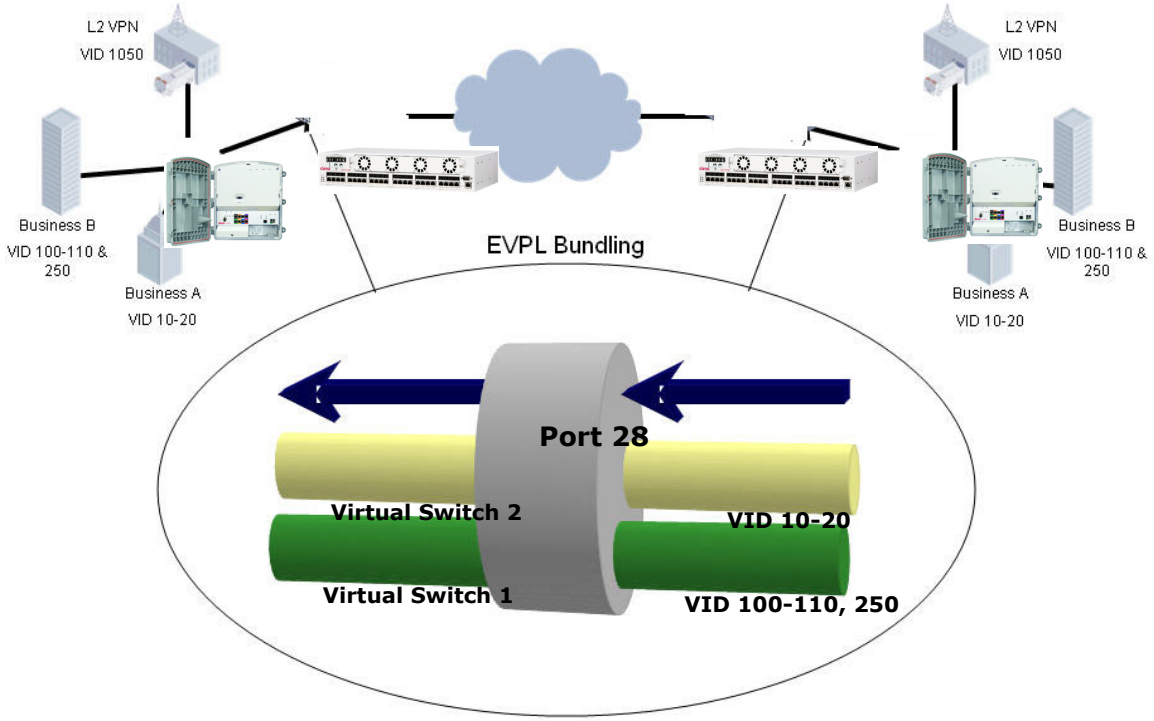

Example

```
> 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.

CVID range configuration

CVID range support is subset of CVID bundling. A CVID range configuration allows a continuous range of VIDs for a given UNI port to be treated as single UNI membership.

In releases prior to 6.20, even if continuous VLAN ids were assigned while adding an EVPL membership of a VS, these VIDs were considered as individual memberships of the VS for a given UNI port.

This configuration caused devices to consume H/W resources for each CVID. Some of the devices (390x and 392x) support only 2048 EVPL members due to hardware resource limitations.

Copyright© 2022 Ciena® Corporation

A CVID range that has more than 2048 VIDs cannot be supported. Additionally, once these resources are completely consumed, it prevents any further configuration of EVPL members on any given UNI port.

CVID Range

The CVID range (for SAOS 6.x devices) supports eight non-overlapping VLAN ranges using hardware range resources for each port. The following devices are supported: 390x, 392x, 3942, 5142, 5160.

The user can configure the CVID range as a single EVPL, which reduces hardware resource usage. Devices that have limited EVPL resources can support CVID ranges with more than 2048 VIDs and still retain enough resources for other EVPL members.

The vlan-range-type specifies the following options for devices using the CVID range.

- Strict: The device creates a single EVPL for the given range, given that both vlan-translation and range resources are available.
- Group: The device creates EVPLs for each CVID present in the given range. This behavior is similar to creating EVPLs without using the vlan-range-type option.
  - The device does not allow the user to remove any individual CVID from the given range.
  - The group option is only available when adding UNI memberships to a Q-in-Q VS where the virtual-circuit has the vlan-transform option set to i-push,e-pop.
  - The group option allows the device to optimize the vlan-translation resources used for the NNI port to achieve i-push,e-pop operation. Optimization depends on the VLAN range configured on the NNI port.

  Note: When the vlan-range-type as group is used, the device may take time to perform add/remove/set/unset operations.

Even though the VLAN range can be configured on any given port, it should only be configured for the NNI ports of a Q-in-Q virtual switch having virtual circuits with vlan-transform set to i-push, e-pop. See Virtual switch SVLAN transforms.

See the following procedures for CVID range support:

- “Configuring the CVID range as a single EVPL”
- “Setting the VLAN range for a port”
- “Unsetting the VLAN range for a port”

Copyright© 2022 Ciena® Corporation

- "Displaying the VLAN range for a device"

Feature interactions

The table below lists the feature interactions for CVID range.

Table 7-12
CVID range feature interactions

Feature	Description
Applicable for both Q-in-Q and MPLS based Virtual Switches	
CFM/Y.1731	UP/Down MEP and Y.1731 sessions are only supported for the first CVID present in the range when the vlan-range-type is set to Strict. The device may not provide a proper error if sessions are created on any VID except first VID in the range.
Benchmark	Benchmark sessions are only supported for individual CVIDs present in the range. A single session for a complete range is not supported.
Access-list	Access-list rules are only supported for the individual CVIDs present in the range.
Q-in-Q Virtual Switch	vs-l2-transform – vlan-range-type is only supported for those UNI ports that have vs-l2-transform set to either i-push,e-pop or i-stamp:push,e-match-pop:stamp. – A UNI port with vs-l2-transform set to i-stamp:push,e-match-pop:stamp cannot use CVID translation for EVPL members that have vlan-range-type as group or strict. – A UNI port with vs-l2-transform set to i-stamp:push,e-match-pop:stamp cannot have any EVPL members with translate CVID set to 0 if they have any EVPL members with vlan-range-type as group or strict. – The device will not match the CVID at the UNI egress when the EVPL is created with the vlan-range-type set to Strict and the UNI port has vs-l2-transform set to i-stamp:push,e-match-pop:stamp. vs-egress-filter – An EVPL membership with the vlan-range-type set as Strict is not supported when the UNI port has vs-egress-filter set to on. UNI-member-statistics – The device does not support statistics for EVPLs having the vlan-range-type attribute set to Strict or Group.

Copyright© 2022 Ciena® Corporation

Virtual switch SVLAN transforms

The 390x, 392x, 3942, 5142, and 5160 platforms support configurable SVLAN transforms for Q-in-Q.

This functionality applies to the NNI at ingress and egress:

- NNI Ingress: classify a CVID and push an sTag
- NNI Egress: ability to pop the sTag

The example in the SVLAN transform: i-push,e-pop for Q-in-Q figure illustrates the following in a Q-in-Q SVLAN configuration:

1. At NNI ingress, the device looks up the frame's VID, which becomes the CVID and classifies it to a virtual switch.
2. Then it pushes the sTag onto the frame so that the frame can be switched in the virtual switch domain internally.

Figure 7-8
SVLAN transform: i-push,e-pop for Q-in-Q

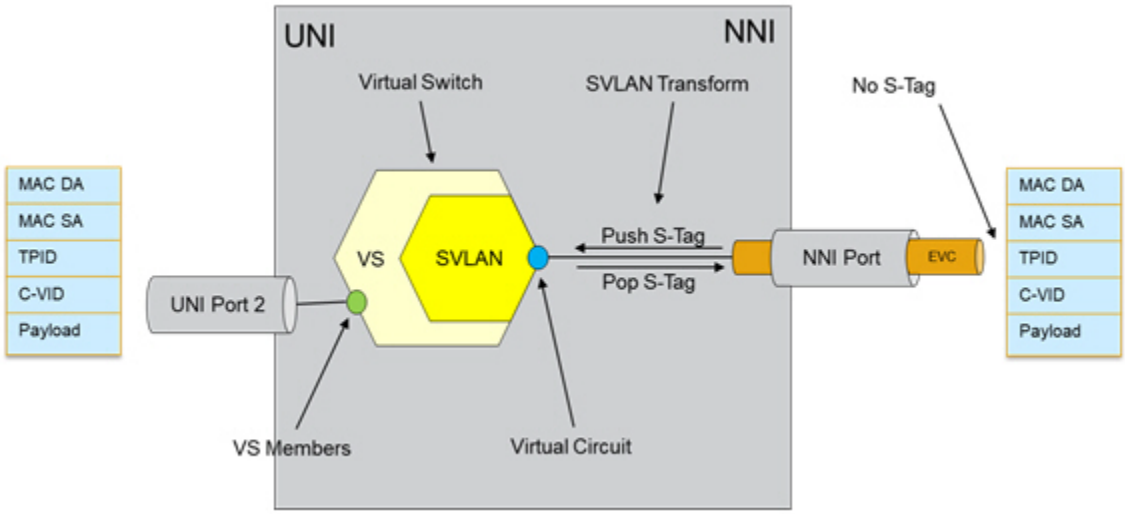

At NNI egress, the devices looks up the frame's SVID only (no CVID). On matching, the device pops the sTag. The frame then egresses out of the NNI as it ingressed on the UNI port without any additional sTag.

The user can configure SVLAN transform using the Vlan transform action i-push, e-pop.

```
>virtual-circuit ethernet create vc <> vlan <> vlan-transform <i-push,e-pop>
```

For Example:

Copyright© 2022 Ciena® Corporation

```
>virtual-circuit ethernet create vc transparent_vc vlan 1000 vlan-transform
i-push,e-pop
```

A device may receive a frame having the same outer vid as SVID. The following is an example of the CLI command to drop such a frame at NNI ingress if needed:

```
>vlan translate add port <nni-port> vid <s-vid> vlan 4095.
```

Note 1: If the CVID range-based EVPL is not required, it is recommended to use CVID translation using the Translate Tag 0 to achieve similar behavior.

Note 2: SVLAN transform is only helpful for transparently switching data traffic. Features that are dependent on host-bound traffic (control plane) are not supported in SAOS 6.20.0 release.

Note 3: In the cases where multiple CVIDs are added using a CVID range configuration, the following CLI can be used to optimize the H/W resources required for NNI port.

```
> port set port <nni-port> vlan-range-index <> vlan-range <>
```

Limitations of SVLAN transform

The following limitations are only applicable to a Q-in-Q virtual switch having a virtual circuit with the vlan-transform set to i-push,e-pop.

- Only a single UNI port can have membership in this type of virtual switch.
- Only a single NNI port can have membership of SVLAN in this type of virtual switch.
- A CVID added as an EVPL in a Q-in-Q virtual switch MUST not match with any of the NNI port's VLAN memberships, including the P-Tag.
  - The CVID can be the same as the SVLAN used for the Q-in-Q VS where this EVPL membership is being added.
  - Since a port MUST have at least one VLAN membership, adding a port as a NNI port may conflict with the CVID of an EVPL. In this case, if the port has a membership of the default VLAN, then the device will remove this membership while adding the port to the SVLAN.
  - Adding an EVPL with a range may conflict with the default VLAN membership of NNI port. In this case, the device will remove the NNI port from the default VLAN.
- A CVID added as an EVPL in a Q-in-Q virtual switch MUST not match with the outer VID of the VLAN translate configured on the NNI port.
- VLAN translate must not be configured on the device that translates any VID to a VID that is also being used as an SVLAN for this type of Q-in-Q virtual switch.

Copyright© 2022 Ciena® Corporation

- EPL memberships are not supported for these types of Q-in-Q virtual switches.
- A port with vs-l2-transform set as i-push,e-pop or i-stamp:push,e-match-pop:stamp can only have UNI memberships of such Q-in-Q virtual switches. CVID translation-related configuration is not supported when a port with vs-l2-transform mode as i-stamp:push,e-match-pop:stamp is added as UNI port in such a Q-in-Q virtual switch.
- A UNI port cannot have untagged-data-vs and/or untagged-ctrl-vs set in this type of Q-in-Q virtual switch.
- Rx/Tx to/from Host CPU is not supported; thus, the following features will not work:
  - L3 interface
  - CFM/Y.1731
  - L2-CFT
    - Peering of untagged L2CP frames from NNI ports is supported.
  - TWAMP
  - Remote Interface
  - G.8032
  - DHCP and DHCP relay
- The egress-untag-vlan attribute will not take effect for NNI ports that have memberships of VLAN that is used to create virtual circuit with vlan-transform i-push,e-pop and egress-untag-vlan is set to same value as SVLAN.
- For CVIDs and SVLAN, vlan-ethertype and vlan-ethertype-policy does not take effect for NNI ports that have memberships of VLAN that are used to create virtual circuit with vlan-transform i-push,e-pop.
- VLAN translation to stamp SVLAN at NNI ingress/egress is not supported.

See the following list of procedures to configure VLAN transform:

- “Configuring a Q-in-Q virtual switch with SVLAN transform”
- “Displaying information related to Q-in-Q virtual switch with SVLAN transform”

Virtual switch L2 transforms

The 3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928, 3942, 5142, and 5160 platforms support configurable VLAN translation to classify Q-in-Q traffic from one L2 switching domain and remap it to another.

Copyright© 2022 Ciena® Corporation

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)"
- "Ingress push; egress pop and stamp (i-push,e-pop:stamp)"
- "Ingress stamp; egress match, pop, and stamp (i-stamp:push,e-match-pop:stamp)"

Ingress push; egress pop (i-push,e-pop)

On ingress, the device looks up the frame's VID (which becomes the CVID) and classifies it 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.

When a UNI port is configured with this L2 transform mode, these 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 do not support translate tag values.

If you wish to do VLAN translations on VS members, the port setting vs-l2-transform must be changed to either i-push,e-pop:stamp or i-stamp:push,e-match-pop:stamp. Untagged frames are supported in i-stamp:push,e-match-pop:stamp mode.

> ***Note:*** EVPL VS members share internal resources with MPLS attachment circuits, VLAN translation and sub-port instances. This means that the scale of these features may be affected by the number of MPLS attachment circuit instances on the switch.

Ingress push; egress pop and stamp (i-push,e-pop:stamp)

On 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.

Copyright© 2022 Ciena® Corporation

When a UNI port is configured with this L2 transform mode, these 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 do not support translate tag values.
- The associated virtual circuit cannot have a VC transform service.

Ingress stamp; egress match, pop, and stamp (i-stamp:push,e-match-pop:stamp)

Currently, an Ethernet (QiQ) VS member and a VLAN member cannot co-exist on the same physical or aggregate port.

A single UNI or customer-facing port may require these services at the same time:

- QinQ encapsulation, where an additional QiQ VLAN tag is added to the ingress frame before being switched out the NNI
- VLAN switching, where frames are simply switched based on the VLAN tag's VID value with or without translation
  - — A common example of VLAN-tagged frames to be switched in this case are those required for device management purposes.

This VS-l2-transform mode provides the ability to support these hybrid services on a single UNI port by supporting two behaviors that differ based on the translate-tag value.

The VS-transform has an ingress and egress action.

Copyright© 2022 Ciena® Corporation

This table describes the ingress and egress UNI behaviors when the translate-tag is non-zero.

Table 7-13
Translate-tag non zero

Translate-tag non-zero	Description
Ingress: i-stamp:push	• The UNI port's ingress action for this transform is first to stamp the incoming customer VID (CVID) with a new VID value, CVID*. The customer VLAN tag (CTAG) has now been translated to CTAG*. The next action is to push on a new service provider tag (STAG). The resulting frame which gets switched by the device has two VLAN tags, the translated customer tag (CTAG*) and the service provider tag (STAG). ***Note:*** The PCP and DEI bits of the customer VLAN tag are not modified.
Egress: e-match-pop:stamp	• The UNI port's egress action for this transform is to first "match" on a subset of a packets. These are packets that the UNI forwards towards the NNI port. These frames have two VLAN tags, the translated customer tag (CTAG*) and a service provider tag (STAG). For only these packets, the STAG popped. Then the remaining VLAN tag, CTAG*, is stamped back to the customer's original CVID. ***Note:*** The PCP and DEI bits of the customer VLAN tag are not modified during this egress transformation.

This table describes the ingress and egress UNI behaviors when the translate-tag is zero.

Table 7-14
Translate-tag zero

Translate-tag zero	Description
Ingress: i-stamp:push	• The UNI port's ingress action for this transform stamps the incoming CVID with a new VID value, SVID. The CTAG has now been translated to a service provider VLAN tag, STAG. The resulting frame gets switched by the device with a single VLAN tag.
Egress: e-match-pop:stamp	• The UNI port's egress action for this transform is to first "match" on a subset of a packets. These frames have one VLAN tag, the STAG. For only these packets, the STAG VID goes back to the original CVID.

This table describes the information that is relevant to the vs-l2-transform modes.

Copyright© 2022 Ciena® Corporation

Table 7-15
VS-l2-transform mode information

Task	Description
Set the VS transform mode of the UNI port to i-stamp:push,e-match-pop:stamp.	This allows VS members on the UNI port to use the translate-tag capability including the translate-tag zero function, which replaces the CVID of the frame with a S-VID on ingress, and the opposite action when egressing the VS member.
Set the traffic to be QiQ encapsulated	• Create a VLAN representing the SVLAN and add NNI ports as members. • Create an Ethernet VC using the SVLAN. • Create an Ethernet VS using the Ethernet VC. • Create a VS member on the UNI port. If an EPL is required, create the EPL as normal. If an EVPL is required, create the EVPL with a translate-tag value equal to the CVID you are specifying for EVPL. This behavior for the QinQ virtual switch is the same as the default VS transform mode of i-push:e-pop for both EPLs and EVPLs since the translate-tag of the EVPL is set to the same value as the CVID, which stamps the CVID with the same value it already has. ***Note:*** EVPL VS members share internal resources with MPLS attachment circuits, VLAN translation and sub-port instances. This means that the scale of these features may be affected by the number of MPLS attachment circuit instances on the switch.

Copyright© 2022 Ciena® Corporation

Table 7-15
VS-l2-transform mode information

Task	Description
Set the traffic to be switched based on VID	• Create a VLAN representing the VID to the switched traffic and add NNI ports as members. • Create an Ethernet VC using the VLAN that was created. • Create an Ethernet VS using the Ethernet VC. • Create an EVPL VS member on the UNI port with the vlan/CVID equal to the VID of the traffic you are switching and a translate-tag of zero. By specifying an EVPL VS member with a translate-tag of zero, this stamps the incoming VID of the frame (CVID) with the VID of the VC (S-VID), and on egress stamp, the incoming VID of the VC (S-VID) with the VID of the EVPL (CVID). Since the EVPL VID (CVID) and VLAN of the VC (S-VID) are the same, there is no change to the VID of the frame and it appears to be a VLAN switched based on the VID. ***Note:*** EVPL VS members share internal resources with MPLS attachment circuits, VLAN translation and sub-port instances. This means that the scale of these features may be affected by the number of MPLS attachment circuit instances on the switch.

There is a single match classification for every VS member added to a UNI port in the "i-stamp:push,e-match-pop:stamp" vs-l2-transform mode. If an egress UNI frame does not match the egress classification, that packet is dropped.

For EVPL and non-zero translate tags, the match occurs on the S-Tag VID and the translated C-Tag VID. For EPL the match occurs on the S-Tag VID. For EVPL with zero translation tag the match occurs on the translated S-TAG VID.

For a virtual switch in this transform mode:

- The translate-tag value must be unique for each EVPL member of a virtual switch on this UNI port.
- The untagged-data-vs can be specified if no EVPL has been added with a translate-tag of zero.
- L2CFT can only be enabled if its associated L2CFT profile is in transparent tunnel mode.
- You cannot have an EVPL with a translate-tag of zero and an EPL on a virtual switch.

Copyright© 2022 Ciena® Corporation

Untagged broadcast, unknown unicast and multicast traffic enhancements

Prior to SAOS 6.15, Private Forwarding Groups (PFGs) were used to prevent broadcast, unknown unicast and multicast (BUM) traffic ingressing UNIs to be flooded to other UNIs, and ensure that it can only egress on NNI. BUM traffic ingressing the NNI destined to a single UNI port was flooded to all UNIs. This raised security concerns. As a result, behavior implemented in SAOS 6.15 includes:

- The UNI ports support and accept untagged frames and associate them to the virtual switch.
- Untagged frames Rx at UNI are given a designated data VID.
- Policy-based and port-based PFG is configured to ensure that UNI-to-UNI traffic is not forwarded.
- BUM traffic that is double tagged (SVID+CVID) has the SVAN tag removed by the virtual switch. The CVLAN tag is also removed by the port, and an untagged frame is presented out the UNI.

To implement this behavior, the port attribute vs-egress-filter option is used. This option only works with vs-l2-transform i-stamp:push,e-match-pop:stamp mode for a UNI port.

Each member has an ingress and/or egress VLAN translation entry associated with a hardware data path. If a traffic flow is not hitting a vlan translation entry, it is discarded.

The multiple types of VLAN translation entries from a NNI towards a UNI are:

- For EVPL, a VS member on a port with enhanced mode has a VLAN translation entry with SVID/CVID as the key for lookup.
- For untagged traffic on the port with untagged-data-vid and untagged-data-vs configured, it has another entry with SVLAN/untagged-data-vid as the key.
- If the whole port is added as an EPL, it has another entry with SVLAN on top to serve as the flooding domain to other traffic flows unspecified in this EPL. All BUM traffic hits this entry.

If the vs-egress filter attribute is turned on, the last entry is not configured in the data path and causes the BUM traffic to drop.

- This option is configured under each UNI port attached to a virtual switch. For BUM traffic from NNI towards UNI direction, traffic can be divided into two types: BUM traffic for a given UNI port identified by destination UNI port ID, and SVLAN, configured CVID or untagged-data-VID. This traffic is allowed for that UNI port only, and it is dropped for other UNI ports if they have the vs-egress-filter option enabled.

Copyright© 2022 Ciena® Corporation

- BUM traffic with SVLAN only or SVLAN with an unknown CVID. This traffic is dropped if all UNI ports have vs-egress-filter enabled.

If a UNI port does not have this option enabled, it can receive either type of NNI BUM traffic.

See the procedure "Configuring flood traffic from a NNI port toward a UNI port" on page 7-91.

Virtual circuit transforms

To configure a VC transform service, you need a virtual switch and a virtual circuit.

The VC transform is applied to a NNI's port of a Q-in-Q service, and identifies enhanced Q-in-Q service. VC transform services enable virtual switches with the capability to add two VLAN tags onto customer traffic flows as they enter the provider network and allows the virtual switch to strip off two VLAN tags from network traffic flows as they enter the customer network.

> ***Note:*** VC transform is supported on 3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928, 3942, and 5142 switches. VC transform is supported on the 5160 by taking a front panel port out of service and assigning the system resources associated with the port to VC transform.

Provider bridging or Q-in-Q, is a networking standard using VLAN tag stacking to support customer traffic tunneling through a service provider's network. Provider bridging uses a second VLAN tag for ingress customer UNI traffic. This second VLAN tag distinguishes the customer's VLAN forwarding domains from the provider's internal forwarding domains. The service provider's tag is the S-Tag, and the customer tag is the C-Tag.

This figure shows the provider bridging frame format.

Figure 7-9
Provider bridging frame format

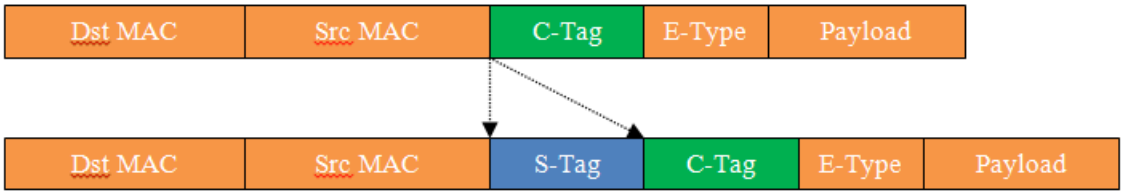

VLAN stacking may apply when more than one VLAN tag is inserted into data frames as they are switched through the provider network(s). VC transforms enhance the existing enhanced Q-in-Q service to push two VLAN tags onto ingress customer flows.

Copyright© 2022 Ciena® Corporation

This figure shows the VC transform frame format.

Figure 7-10
VC transform frame format

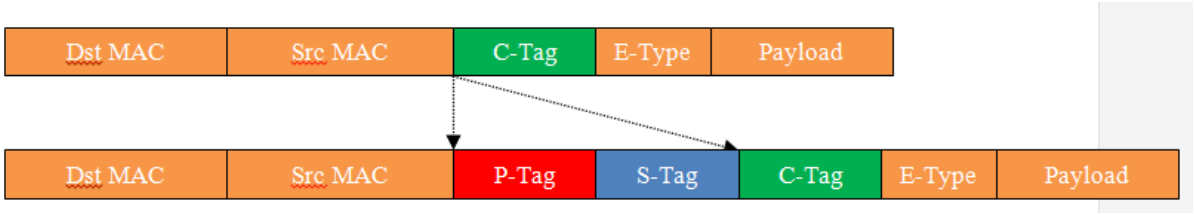

This figure shows an example of a customer site that spans two carrier Ethernet networks.

Figure 7-11
VC transform example

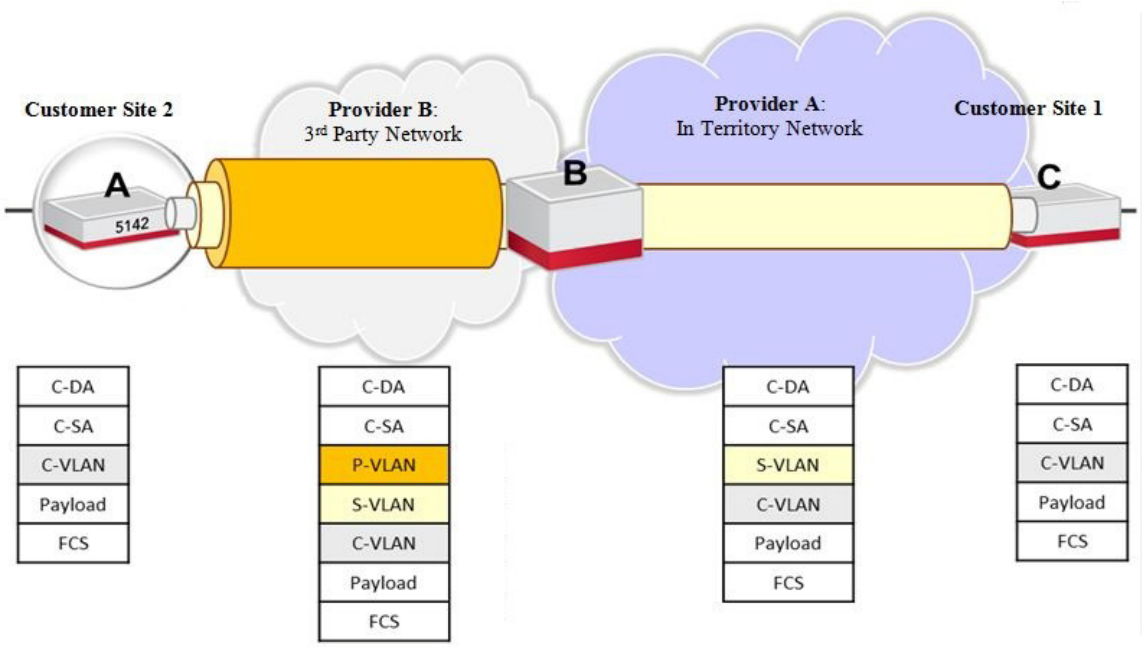

Before traffic enters the third-party carrier Ethernet network, Provider A's switch (node A) must insert two VLAN tags on customer traffic. The first VLAN tag identifies the service in Provider A's network (S-Tag). The second VLAN tag identifies the service in the third-party network. All services are switched in the third-party network using the same VLAN ID, meaning that the P-Tag frames may have various S-Tags going into Provider B's network.

At the "in territory" network, a single tag is stripped off and packets are switched towards node C. At node C, these packets are Q-in-Q decapsulated to expose the underlying customer frame.

Copyright© 2022 Ciena® Corporation

In this example, Provider A is the primary service provider for a customer, but Provider A does not service the new customer site. Provider A uses Provider B, a third-party provider. Provider B services the new area as a bridge between customer endpoints. A VC transform at node A pushes two provider VLAN tags onto customer frames entering the UNI and pops off two VLAN tags for the encapsulated customer frame entering the NNI. The first VLAN tag is the VLAN tag for Provider A called the S-Tag. The second VLAN tag is the VLAN tag for Provider B called the P-Tag.

> ***Note:*** At node A, all switching decisions are based on the S-Tag. The P-Tag is pushed on egress of the NNI port does not determine the internal forwarding domains or resolved CoS values.

This figure shows an example of a virtual switch with a VC transform at node A.

Figure 7-12
VC transform at node A example

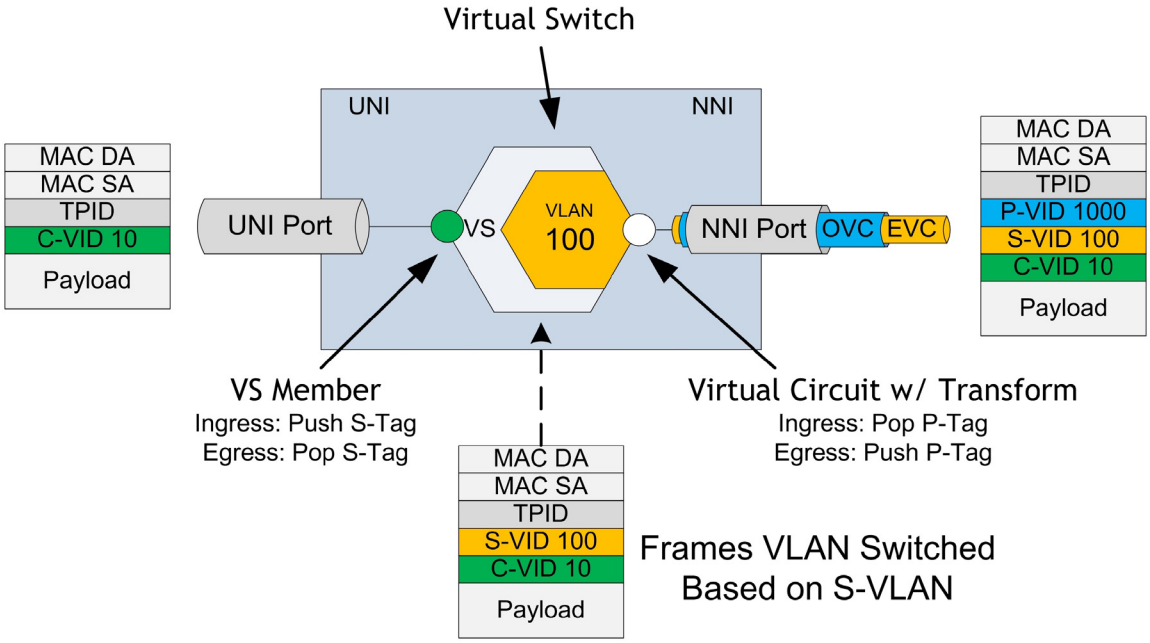

Configuration model

VC transforms plug directly into the existing Q-in-Q virtual-switch and VS cross-connect configuration models.

VC transforms are created and configured as individual objects in the user interface. These objects are used as containers for additional VLAN tag encapsulation information which is applied to packets egressing all associated

Copyright© 2022 Ciena® Corporation

virtual circuits. Several VC transforms can be configured and used on various Q-in-Q VCs to provide new VC transform services or enhanced Q-in-Q services.

Q-in-Q virtual switch configuration model

This figure shows the relationship between a VC transform and a virtual circuit.

Figure 7-13
VC transform and virtual circuit relationship with Q-in-Q VS

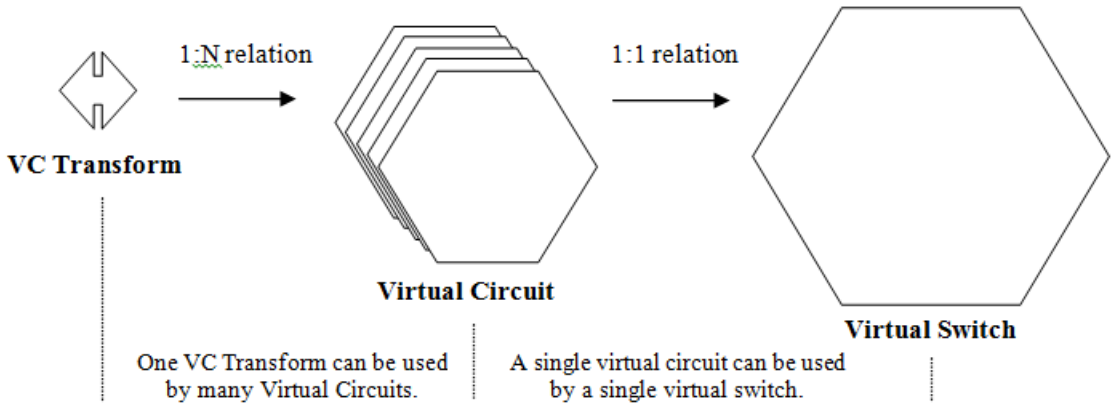

Configuring a VC transform service is similar to configuring a Q-in-Q service. A virtual circuit and virtual switch are configured to define the logical connection between one or more customer endpoints. For Ethernet Q-in-Q services, a virtual circuit specifies a configured VLAN on the device as the SVLAN ID used to encapsulate traffic from the UNI and decapsulates traffic from the NNI. This VLAN is also used as the forwarding domain on the switch. MAC addresses are learned against this VLAN.

You can specify an L2 transform to enable double-tagged encapsulation of UNI traffic. A VC with an L2 transform specified identifies a VC transform service. For ingress UNI traffic, this service performs Q-in-Q encapsulation and pushes on an extra (second) 802.1Q VLAN tag before egressing the virtual circuit's NNIs ports. In the opposite direction, the packet enters the NNI and its outer VLAN tag is stripped off. This packet is then switched based on the remaining S-Tag, and before leaving the UNI port, the S-Tag is removed. Learning and forwarding are based on the virtual circuit VLAN of the service.

A VC which identifies an L2 transform can be associated with a virtual switches to provide enhanced Q-in-Q functionality. The SVLAN specified during VC creation defines the NNI ports for the service while the associated virtual-switch EPL and EVPL members define the UNI members.

Copyright© 2022 Ciena® Corporation

VS cross-connect configuration model

The VS cross-connect configuration model is similar to the model for configuring a VC transform service for a Q-in-Q service. To apply this service, an Ethernet virtual circuit that is associated with one or more VS cross-connects is configured with a VC transform.

This model lets a device switch Time Division Multiplexing (TDM) frames on a VS cross-connect using a TDM pseudo wire to process the frames. The TDM pseudo wire then transmits or receives frames by an Ethernet virtual circuit with an extra VLAN tag.

This figure shows the VS cross-connect VC cross-connect configuration model.

Figure 7-14
VS cross-connect VC cross-connect

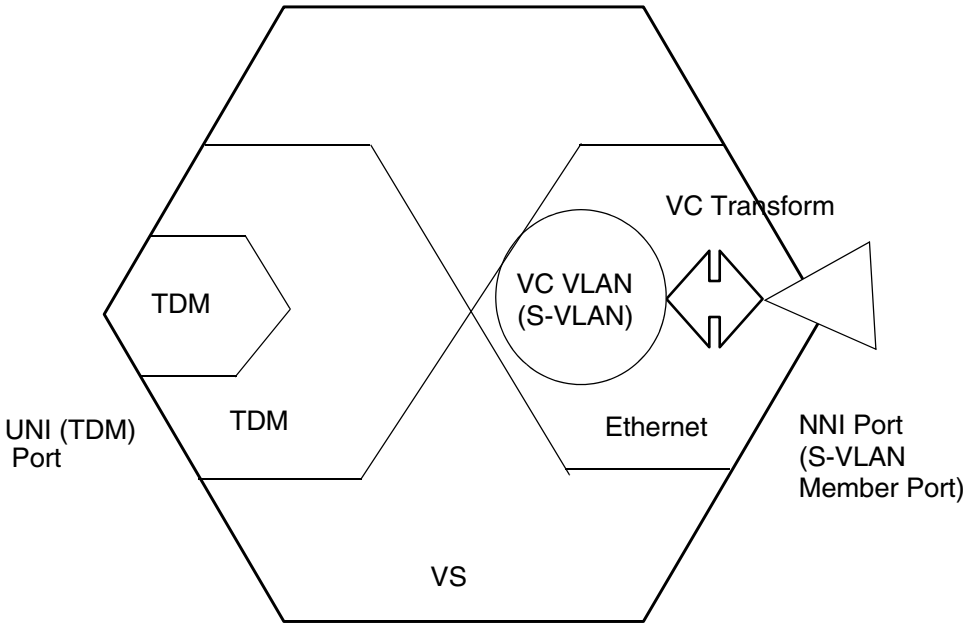

VC transform creation

A virtual circuit for a VC transform service is different than one used on a Q-in-Q service because of the additional L2 transform that is applied to the virtual circuit. An L2 transform object is supported under the existing virtual-circuit object. This L2 transform can be applied to a VC on creation or after creation using the set command.

Copyright© 2022 Ciena® Corporation

A virtual circuit transform is defined by various VLAN tag attributes, as described in this table.

Table 7-16
VLAN tag attributes

VLAN tag attribute	Description
vid	The VLAN ID that sets in the VLAN tag is pushed on egress, and also used as ingress match criteria for this VC. For example, if the associated VC uses an SVLAN ID of 100 and the specified transform pushes a VLAN ID of 888 on egress, this VC matches on all ingress frames with an outer tag of 888 and an inner tag of 100.
pcp	The priority code point (pcp) value set in the VLAN tag that is pushed on at egress. The pcp attribute is optional and defaults to 0 if not specified.
dei	The drop eligible (dei) value set in the VLAN tag that is pushed on at egress. The dei attribute is optional and defaults to 0 if not specified.
tpid	The tag protocol identifier (tpid) value set in the VLAN tag that is pushed on at egress. The tpid attribute is optional and defaults to 0x8100 if not specified.

When a VC transform object is created it can be referenced by many virtual circuits. This prevents duplicating the same action and tag attributes for various virtual switches or VS cross-connects that perform the same transform. Virtual switches or VS cross-connects which push on the same P-tag can reference the same VC transform object.

A NNI port cannot be used to push on different P-Tag values. A NNI port must be associated with the same VC transform. All virtual circuits which use a particular NNI port must reference the same VC transform.

Copyright© 2022 Ciena® Corporation

The maximum number of transforms is limited to one. The maximum number of virtual circuits that support transforms is shown in this table.

Table 7-17
Maximum supported VCs

Platform	Maximum supported VCs with VC transforms	Maximum NNI ports participating in VC transform service
3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928, 3942	512	1
5142, 5160	1024	1

Applying a VC transform to a VS

A VC which contains an L2 transform is configured with a virtual switch the same way as a VC which does not contain an L2 transform. This VC is associated to a virtual switch on creation or after a VS is created using the set command.

Applying a VC transform to a VS cross-connect

A VC which contains an L2 transform is configured with a virtual switch cross-connect the same way as a VC which does not contain an L2 transform. This VC is associated to a virtual switch cross-connect only on the creation of the cross-connect which also includes the TDM VC.

CFM on a VC transform service configuration

Support of CFM and Y.1731 on VC transform services is limited. See "VC transform feature interactions" on page 7-41.

The configuration for CFM/Y.1731 based-services on UNI ports is the same as Q-in-Q-based configuration on UNI ports. You can enable CFM on a VC transform VS and configure both up and down MEPs. Frames originating from an up MEP UNI egress the corresponding NNI ports with at least two VLAN tags. The outer most tag is the P-Tag, pushed on by the VC transform, while the inner VLAN tag is the S-Tag, representing the VC VLAN. Frames originating from a down MEP UNI egress the UNI port with the same tag behaviour seen in equivalent Q-in-Q configurations.

With respect to configuring CFM/Y.1731-based services on NNI ports for VC transform services, a CFM up or down MEP is not permitted on NNI port members. Instead, the NNI link, represented by P-Tag tagged frames, is monitored by creating a VLAN which matches the target P-Tag service VLAN ID. The NNI port is then added to the newly created VLAN and CFM is configured against this VLAN-based service. CFM down MEPs configured on the NNI port egress the switch with the P-Tag service VLAN ID.

Copyright© 2022 Ciena® Corporation

Resource manager

VC transforms require the allocation of hardware resources. By default, no resources are assigned. This resource usage applies:

- For 3903, 3903x, 3904, 3905, and 3906 switches, one resource block supports up to 256 virtual switches or VS cross-connects with Ethernet virtual circuits that use VC transforms. A total of 160 static classifier resources are consumed on allocation of the first VC transform resource block.
- For 3924, 3926, 3928 and 3942 switches, one resource blocks supports up to 256 VSs with VC transforms. A total of 158 static classifier resources are consumed on the allocation of the first VC transform resource block.
- For 5142 switches, one resource block supports up to 512 virtual switches or VS cross-connects with virtual circuits that use VC transforms. A total of 286 static classifier resources are consumed on the allocation of the first VC transform resource block. A maximum of 1024 VSs or VS cross-connects support VC transforms.

VC transform feature interactions

This table describes the feature interactions with VC transform services.

Table 7-18
VS transform feature interactions

Feature	Description
L2 control frame tunneling (L2-CFT)	Virtual switches which use a VC with L2 transforms only supports transparent tunneling of control frame traffic. An L2 CFT tunnel method of L2PT is not supported on VC transform virtual switches.
MC LAG	MC LAG is not supported on VC transform NNI ports.
MAC learning	Learning is based on the VLAN ID of the S-Tag for VC transform services (the virtual circuit LAN).
G.8032	G.8032 is not supported on VC transform services.
CoS mapping	A VC transform can define a P-Tag with a fixed frame CoS value. The VCs using this VC transform have the same P-Tag attributes. All services configured on a particular NNI port which use VC transforms use the same fixed frame CoS value for the P-Tag. Frames which ingress a VC transform service's NNI port(s) derive an RCoS value based on the PCP value of the S-Tag, not the P-Tag.
Management interface	A VC transform service can be indirectly set as the management interface similar to Q-in-Q VSs. The SVLAN associated with VC transform service can be configured as the management VLAN to support double-tagged management traffic.

Copyright© 2022 Ciena® Corporation

Table 7-18
VS transform feature interactions

Feature	Description
VLAN translation	VLAN translation cannot be configured to translate a VLAN ID which matches a P-VLAN ID that is associated with a VC transform on a virtual circuit's NNI port members. This restriction applies to double-tagged VLAN translation configurations. This means that you cannot configure a double-tagged VLAN translation instance where the outer VLAN ID matches a P-VLAN ID that is associated with a VC transform on a VC's NNI port members.
Port VS L2 transforms	The port VS L2 transform "i-push,e-pop:stamp" is not supported on a port that has a VS member belonging to a VC transform VS. The port VS L2 transform "i-stamp:push,e-match-pop:stamp" is supported on a port that has a VS member belonging to a VC transform VS, but that VS member cannot have a translate-tag with a non-zero value. VS members which do not specify a translate tag do not adhere to the traditional VS L2 transform mode's match behavior which matches on the S-Tag and C-Tag of the egressing UNI frames. Instead, these VS members match on only the S-Tag. The impact is that traffic with non-matching CVIDs can be switched out UNI VS member ports.
Untagged Data VID	The untagged-data-vid port attribute is configurable for UNI ports on a VC transform service. The egress action of this attribute is not supported. This means that a frame that ingresses the NNI port of a VC transform service does not have the CVID stripped off when egressing a UNI port which has the untagged-data-vs set to VC transform service, and the untagged-data-vid attributes set to CVID of the egressing frame.

Copyright© 2022 Ciena® Corporation

Table 7-18
VS transform feature interactions

Feature	Description
CFM/Y.1731	A non-accelerated CFM service can be configured on a VC transform virtual switch. This service only supports MEP creation on VS member ports. Accelerated CFM services cannot be configured on a VC transform virtual switch. Accelerated Y.1731 sessions are supported on a VC transform virtual switch providing that the equivalent Q-in-Q configuration (without a VC transform) is supported. Y.1731 sessions that are triggered using the “frame-loss” CLI are not supported on VC transform services. The VLAN associated with the VC transform P-Tag can be used for either data or CFM monitoring of the NNI link, but not for both. When you configure a VC transform service where the VC VLAN is equal to the VC transform VLAN ID, then they are unable to reuse the same VLAN ID to provide CFM monitoring to the NNI link.
VLAN cross-connects	VLAN cross-connects which share the same outer VLAN ID as the VC transform P-Tag attribute is not supported.
Sub-ports	Sub-ports are supported on the same VC transform NNI port, but ingress VC transform classification always takes precedence over single tagged ingress sub-port classification.

Copyright© 2022 Ciena® Corporation

VC transform considerations

This table describes the VC transform considerations.

Table 7-19
VC transform considerations

Limitation	Description
Supported NNI ports on VC transform service	The number of NNI ports that can be used on a VC transform service is limited to one. This means that a VC which specifies a VC transform can have one VLAN member in the associated SVLAN.
Supported VS UNI ports on a VC transform service	The number of virtual switch UNI ports that can be used on a VC transform service is limited to one. This means that a VS which uses a VC that identifies VC transforms is only able to have one VS member port. This restriction does not limit the number of EVPL members the VS can have. It restricts the number of physical, or aggregation, ports that can participate as a UNI in the target VC transform service.
VC transform on NNI ports	VC VLANs with a VC transform using the same NNI port must reference the same VC transform. This means that a NNI port can be used by multiple VC transform services, but these services must have the same VC transform.
VC transform supported TPIDs for frames received on NNI ports.	Frames received by a NNI port are only identified and switched by a virtual switch with a VC transform service if the TPIDs in the frame contain these values: • P-Tag (outer VLAN tag) TPID: 0x8100, 0x88a8, or 0x9100 • S-Tag (inner VLAN tag) TPID: 0x8100 only
Egress mirroring	Port mirroring is not supported on a VC transform's NNI port. This port must not be used on an ingress or egress mirror.
Mirroring to a VC transform NNI port	A network facing port that uses VC transforms cannot be used as a mirror port. You cannot configure any mirror encapsulation on this port.
Ingress mirroring	A network facing port that uses VC transforms can ingress mirror traffic, but some control frames may not be mirrored.

Copyright© 2022 Ciena® Corporation

Table 7-19
VC transform considerations

Limitation	Description
Accelerated Y.1731 sessions on VC transform NNI ports	When a VC transform service is configured on a particular NNI port, that NNI port does not support accelerated Y.1731 sessions for any service. Accelerated services configured on this port incorrectly inherits the VC transform P-Tag and incorrectly egresses the NNI port with an extra VLAN tag. When configuring a VC transform service where the VC VLAN and the VC transform P-tag are set to the same value, the VC VLAN does not support CFM services for both the transform service and the NNI link. This means that the VLAN associated with the VC transform P-Tag can be used for either data or CFM monitoring of the NNI link, but not for both.
Unsupported queues on VC transform ports	All ports, UNI or NNI, involved on a VC transform service do not support using queue 7. This means that any VC transform service port must not use a queue mapping where an RCoS value is mapped to queue 7. ***Note:*** In the default port settings, queue 7 is not mapped.
Hierarchical queuing	Hierarchical queuing, also known as service-based queuing, is not supported on VC transform services.

See these procedures:

- "Configuring VC transform support on the 3928 and 5160 switches" on page 7-95
- "Configuring a virtual circuit transform service" on page 7-97
- "Configuring a remote management interface on a VC VLAN on a VC transform" on page 7-100
- "Configuring a virtual circuit transform to a virtual circuit" on page 7-103
- "Removing a virtual circuit transform from a virtual circuit" on page 7-104
- "Setting the virtual circuit transform to default values" on page 7-105
- "Enabling or disabling a virtual circuit transform" on page 7-106
- "Deleting a virtual circuit transform" on page 7-107
- "Displaying virtual circuit transform" on page 7-108
- "Preventing NNI ingress frames without P-Tag from being switched" on page 7-111

Private forwarding groups

Layer 2 private forwarding groups (PFG) restrict forwarding of L2 traffic between groups of ports in a shared VLAN or virtual switch. Port-based PFGs alter the forwarding behavior of ports such that their forwarding domain is

Copyright© 2022 Ciena® Corporation

restricted to a subset of ports on the switch. Unlike L2 PFG which restricts the flow of traffic between ports on the same VLAN, port-based PFG is applied to all VLANs. 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 in a particular forwarding domain.

This figure shows an example where the core network connects to the customer's access network using two metro aggregation devices.

Figure 7-15
Multi UNI/NNI example: L2 Network topology

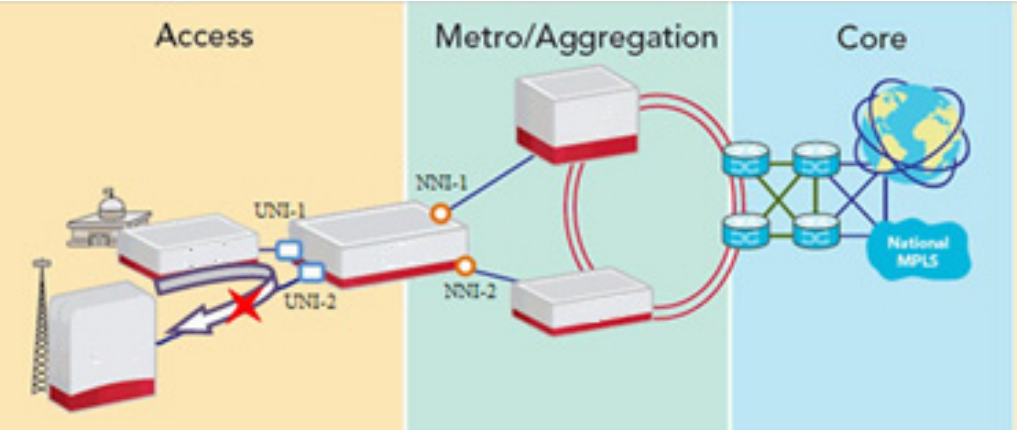

In this configuration, both customer ends are provisioned for core access, but at the same time communication is restricted between customer sites. This can be accomplished by defining forwarding rules for specific groups of interfaces in a forwarding domain. By separating NNIs and UNIs into two forwarding 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.

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 in the same group do not forward traffic to each other.

The following figure shows an example where a port-based PFG facilitates a single switch having more than two sets of private forwarding groups. In this example, the switch provides services to four networks from two core networks. Networks UNI-1 and UNI-2 are connected to core network NNI-1. Networks UNI-3 and UNI-4 are connected to core network NNI-2.

Copyright© 2022 Ciena® Corporation

Figure 7-16
Multi-UNI/Multi-NNI example: Port-based network topology

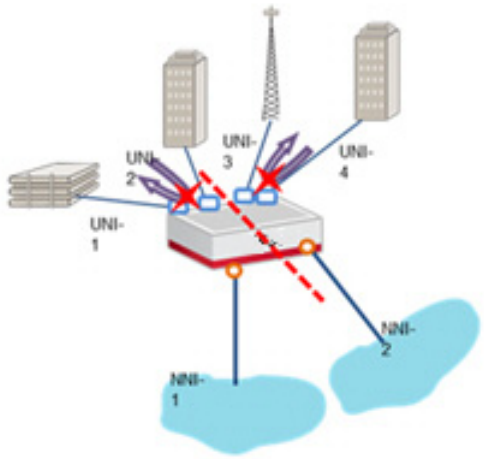

As with L2-based PFG, communication between sites must be restricted. Communication between core networks may be restricted. By restricting which ports may egress traffic received on a given port, forwarding domains can be provided for each port.

L2-based PFG Forwarding 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 in that particular forwarding group. A forwarding policy takes on one of these values:

- A: can only forward to port members of private forwarding group A

Copyright© 2022 Ciena® Corporation

- 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

In the default configuration, all ports in PFG A can forward frames to all other ports on the switch, including ports in 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.

This figure shows the structure of the policy set.

Figure 7-17
Policy set structure

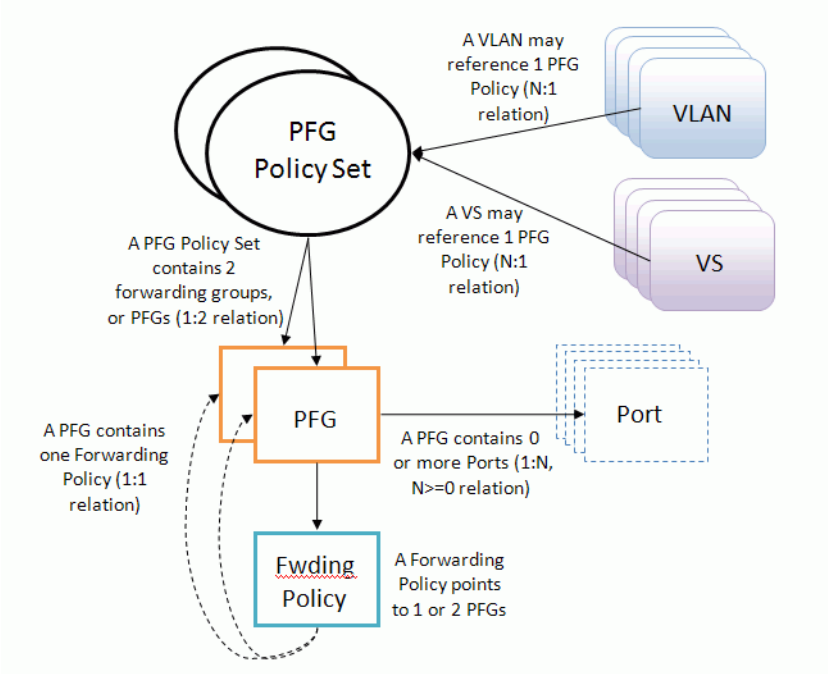

Copyright© 2022 Ciena® Corporation

With L2-based PFG, traffic flow is restricted by placing ports into forwarding groups and defining a set of rules between the groups. For example, PGA only forwards to PFG B, as shown in this figure.

Figure 7-18
L2-based PFG filter

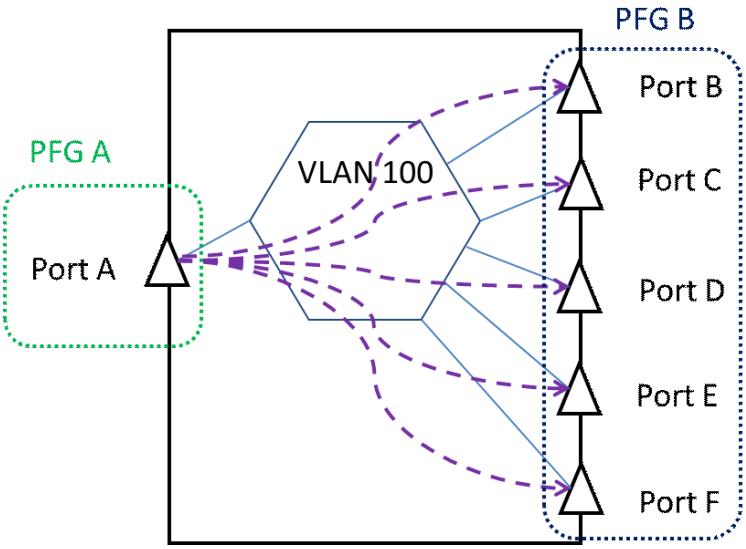

Port-based PFG egress profile

Port-based PFG defines forwarding restrictions for all data traffic on a set of ports which share an egress forwarding policy. An egress forwarding policy defines which ports may egress traffic which has been received on a port to which this policy is applied.

This figure shows a port-based PFG structure.

Copyright© 2022 Ciena® Corporation

Figure 7-19
Port-based PFG structure

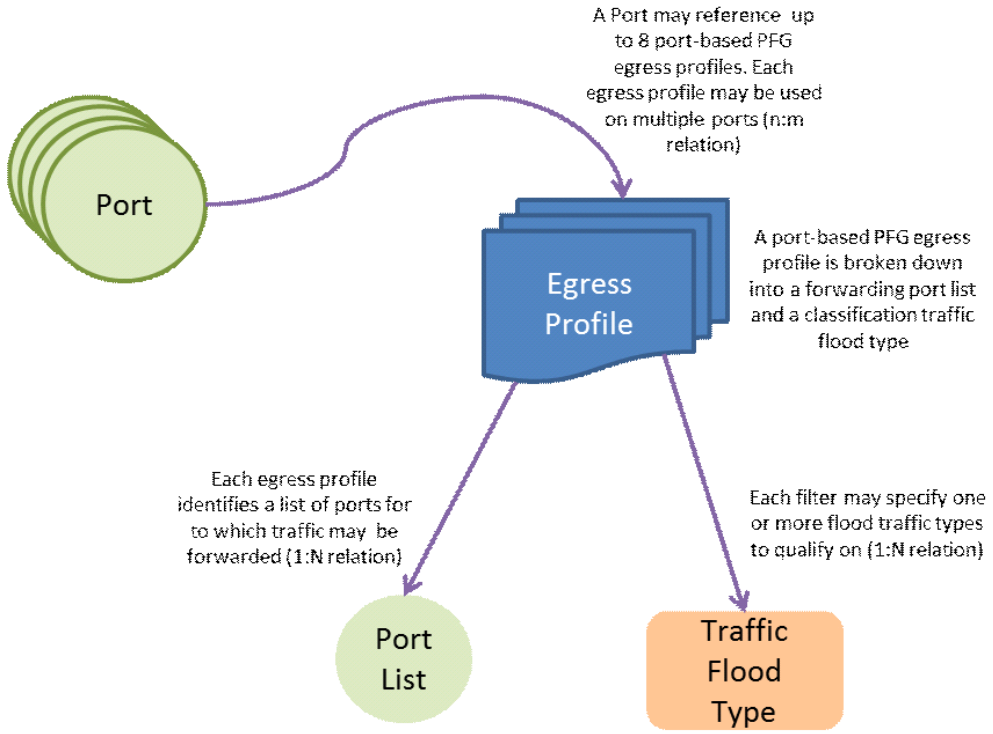

In the example in the Port-based PFG egress profile figure, the egress profile allows traffic to be forwarded to ports B, D and F. The egress profile is assigned to port A. Each egress profile is unidirectional, such that each port must be assigned an egress profile if the data traffic it receives is to be restricted. Bidirectional port forwarding may be configured by creating an egress profile for a set of ports, and then assigning that egress profile to the same set of ports.

Copyright© 2022 Ciena® Corporation

Figure 7-20
Port-based PFG egress profile

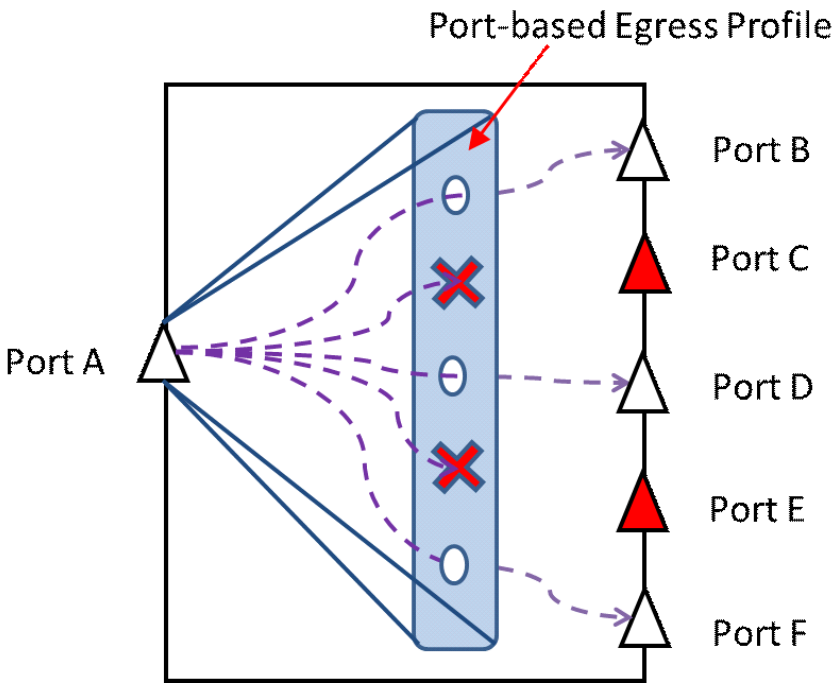

A port-based PFG egress profile defines forwarding restrictions for all data traffic on a particular port. It may be applied to either Link Aggregation Groups (LAGs) or physical ports. The same filter can be used on multiple ports.

Port-based PFG operates co-operatively with L2-based PFG. A port only forwards to another port if port-based and L2-based rules allow it.

Traffic which is switched by the CPU/software is not affected by port-based PFG restrictions. This includes frames which are generated by or destined for the CPU, such as Connectivity Fault Management (CFM). CFM frames which are not switched by the CPU obey the port-based PFG rules. This includes RFC 2544 frames. If an RFC 2544 Port Under Test (PUT) is blocked by port forwarding rules, it does not reflect, analyze or generate test frames. TWAMP test and control PDUs obey port forwarding rules. L2CP frames which are switched by the CPU (Tunneled Traffic for L2PT MAC conversion) do not obey both L2-based and port-based PFG rules.

Copyright© 2022 Ciena® Corporation

Port membership

Physical ports on 39XX/51XX switches are divided into L2-based PFGs on system boot. Each L2-based PFG policy contains its own set of forwarding groups. The L2 PFG membership of a port in a PFG policy is configurable at any time through both the CLI and SNMP interfaces. Port membership also extends to aggregated ports through LAGs.

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 L2 PFG membership (the forwarding group before entering into the LAG) is restored.

This table lists default values for PFG A and PFG B.

Table 7-20
Platform PFG port membership default values

Platform	PFG A	PFG B
3903	1-3	---
3903x	1-3	---
3904	1-4	---
3905	1-4	---
3906	1-3	---
3924	1-8	---
3926	1-8	---
3928	1-12	---
3942	1-24	---
5142	1-24	---
5160	1-24	---

Port-based PFGs are defined similarly for UNI and NNI ports. There are no restrictions on the number or membership of egress profiles.

As with L2-based PFG, port membership extends to aggregated ports through LAGs. Any port that is added to a LAG inherits the egress profile of that aggregation. When a port is removed from the aggregation, the port’s original egress profile membership is restored. A newly created LAG is not a member of or associated with any egress profile.

Copyright© 2022 Ciena® Corporation

Feature interaction

Since the PFG feature alters the forwarding rules of provisioned VLANs and virtual switches, consideration must be given when configuring other features that use the same VLAN service. This table describes the PFG feature interactions.

Table 7-21
Feature interactions

Feature	Description
Port Mirroring	The forwarding rules outlined by a L2-based PFG policy do not affect the ability to mirror traffic on a particular port. L2-based PFG has no effect on port mirroring. Port-based PFG has no effect on ingress port mirroring.
Rapid Spanning Tree Protocol	The spanning tree state of a port always takes precedence over the forwarding rules of a PFG policy. A PFG forwarding policy may state that port 5 on the device is allowed to forward frames to port 8. If port 8 has a STP state of blocking, any packets defined for this port are not forwarded.
G.8032	A VLAN may be protected by a G.8032 and a PFG policy. A warning message is displayed if a user attempts to configure a G.8032 ring on a VLAN with private forwarding groups enabled and forwarding rules have restricted traffic between ring ports. If a user attempts to enable private forwarding groups on a VLAN that is already protected by a G.8032 ring, a warning message is displayed if traffic between ring ports is prevented. A warning message is also displayed if port-based PFG configuration blocks traffic between east and west logical ring ports, or between client ports and ring ports if there are any virtual rings configured for that logical ring. A logical ring with no configured virtual rings does not display any connectivity warning messages if any traffic is blocked between its east and west ports by PFG.
Traffic Profiling	Traffic profile metering is not affected by port-based PFG. All ingress traffic on a port is metered even if some of it is blocked by port forwarding rules. This applies to narrow and wide classifier modes.
CFM	CFM frames which are not switched by the CPU, obey the port-based PFG rules.
RFC 3544	If an RFC 2544 port under test is blocked by port forwarding rules, it does not reflect, analyze or generate test frames.

Copyright© 2022 Ciena® Corporation

Table 7-21
Feature interactions

Feature	Description
TWAMP	TWAMP test and control PDUs obey port forwarding rules.
L2CFT	L2CP frames which are switched by the CPU (tunneled traffic for L2PT MAC conversion) do not obey both L2-based and port-based PFG rules.

Upgrading and downgrading a device

When a device is upgraded from a version of SAOS that does not support L2-based PFG to one that does, most of the default L2-based PFG feature settings are inherited. The policy sets are disabled and the management VLAN is PFG disabled. Usually, when a device boots with default settings, PFG is enabled on the management VLAN. The default PFG port-memberships may conflict with a previously saved configuration. This can prevent the communication of management traffic between provisioned services.

Upgrading a device from a version of SAOS that does not support port-based PFG to one that does has no port-based PFG restrictions.

Downgrading a device from an L2-based or port-based PFG release to a release which does not support PFG requires the removal of all PFG commands from the device's configuration file. This is required to eliminate errors when loading any previously saved configuration files.

External Network-to-Network Interface Hairpin

External Network-to-Network Interface (ENNI) transforms and switches frames out of the same physical port they came in on. Hairpin switching occurs when a frame from an origin endpoint is returned in the direction it came from to get to its destination endpoint. This is possible when an Operator Virtual Connection (OVC) associates two or more OVC endpoints at the ENNI. It allows traffic to be forwarded among multiple service flows of one ENNI port or Link Aggregation Group (LAG).

Improper use of Hairpin Switching can result in a data loop between two operator Carrier Ethernet Networks (CENs) at a single ENNI.

> ***Note:*** ENNI Hairpin is supported on the 3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928, 3942, 5142, and 5160.

Copyright© 2022 Ciena® Corporation

Sub-port interfaces

Sub-port interfaces are required to support ENNI hairpin switching. The number of sub-port interfaces and sub-port-based virtual switches is impacted by the configuration of other features and services in the device. Sub-port interfaces provide these capabilities:

- Create, delete and modify a sub-port interface
- Attach or detach a sub-port interface to or from an Ethernet virtual switch
- Support a sub-port interface based on a port and a single outer VLAN ID value:
  - — Ingress frame classification for the sub-port interface is based on the ingress port and outer VLAN ID value of the received frame.
  - — Frames egressing a tagged sub-port interface have their outer VLAN ID value stamped to the VID configured for the sub-port interface. If the frame is untagged when received on the ingress sub-port, a single VLAN tag is pushed on the frame with the VID of the egress sub-port.
- Support a sub-port interface based on port and untagged frames
  - — Ingress frame classification for the sub-port interface is based on the ingress port and the received frame being untagged. A tagged sub-port only qualifies on untagged or priority tagged ingress traffic.
  - — Frames egressing an untagged sub-port interface have their most outer VLAN tag popped off the frame. If the frame was untagged when received on the ingress sub-port, the frame remains untagged when egressing the egress sub-port.
- Receive and transmit statistics on a per sub-port interface based on frame and byte counts
- Parent port co-existence with VLAN members which is beneficial for
  - — Other non-hairpinning OVC endpoints on the ENNI interface
  - — Remote management on the ENNI-N device and through the ENNI interface

  Note: Sub-port interfaces share internal resources with EVPL VS members, MPLS attachment circuits, MPLS virtual circuits and VLAN translation instances. This means that the scale of these features may be affected by the number of sub-port instances on the switch.

Sub-port configuration restrictions

This list outlines the sub-port configuration restrictions:

- A sub-port cannot be attached to an MPLS virtual switch whose mode is set to vpls or vpws.
- A sub-port cannot co-exist on the same virtual switch with a VS member or a virtual circuit. This means that

Copyright© 2022 Ciena® Corporation

— A sub-port cannot be attached to a virtual switch that currently contains a VS member or a virtual circuit.

— A VS member cannot be added to a virtual switch that has a sub-port attached to it.

— A virtual circuit cannot be set on a virtual switch that has a sub-port attached to it.

- A sub-port cannot co-exist on the same physical or aggregate port with a VS member. This means that

  — A sub-port cannot be created where the parent port is currently part of an existing VS member.

  — A VS member cannot be added where the port specified is a parent port of an existing sub-port.

Sub-port configuration options

When the parent port of the sub-port interface is specified, there are two configuration options on a sub-port interface to further perform classification and egress L2 transforms. These options are

- vtag-stack <vid> — received VLAN tagged frames with the outer VID matching the configured VID are handled by the sub-port interface. Frames being transmitted by the sub-port interface have their most outer VLAN tag stamped with the configured VID by the device before transmission. If the frame is untagged, a VLAN tag with the configured VID is pushed onto the frame by the device before transmission.

  This figure shows tagged sub-port switching behavior.

Figure 7-21
Tagged sub-port switching behavior

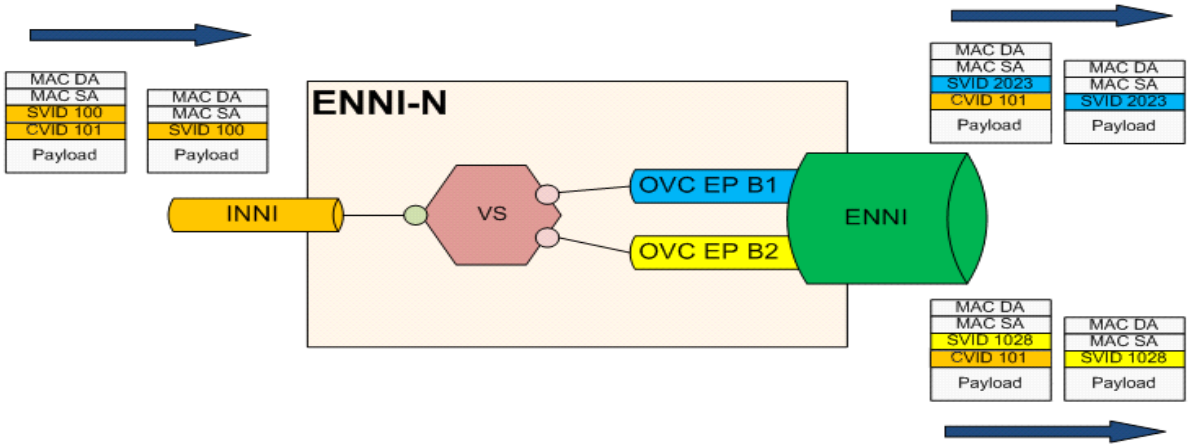

Copyright© 2022 Ciena® Corporation

- vlan-untagged-data — received untagged or priority tagged frames are handled by the sub-port interface. Frames transmitted by the sub-port interface have their most outer VLAN tag popped off by the device before transmission, if one exists.

  This figure shows untagged sub-port switching behavior.

Figure 7-22
Untagged sub-port switching behavior

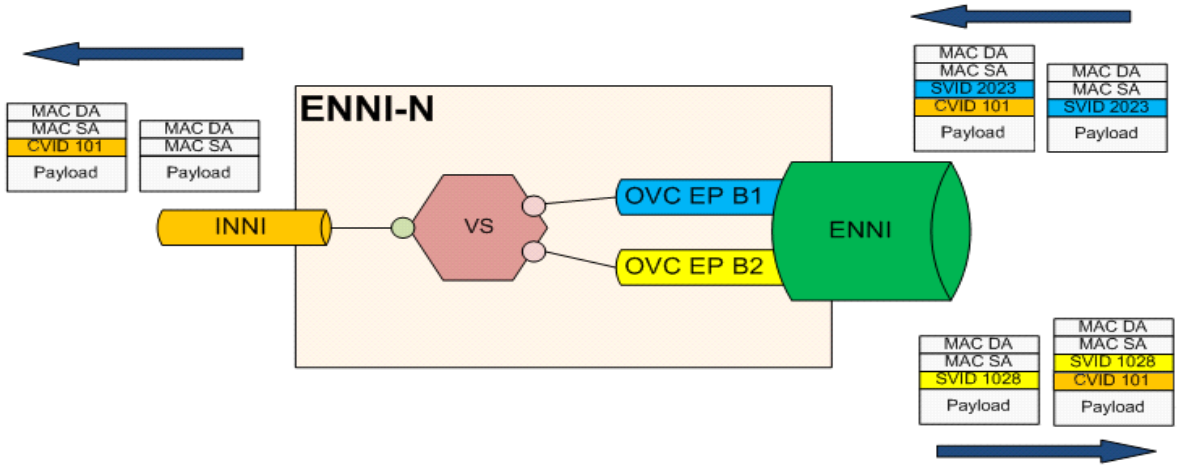

Sub-port statistics

Sub-port statistics and UNI member statistics share egress resources to collect Tx statistics. Both features can consume all the shared resources.

See the procedure "Determining egress resources to collect Tx statistics" on page 7-78 to check resource usage.

Virtual switch behavior

The configuration of an Ethernet virtual switch to attach sub-ports is not affected by ENNI Hairpin Switching. Frames are internally L2 switch-based on the reserved VLAN. The reserved VLAN is never seen on the frame outside the device unless the egress sub-port is configured with the same VID.

For MAC learning, frames received by sub-ports may be learned. If they are learned, MAC learning entries are stored and referenced with this information:

- Key: Virtual Switch, MAC Address
- Destination: Sub-port

Copyright© 2022 Ciena® Corporation

MEF services

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

- MEF 6.2 services: Associating two more UNIs in a service provider Carrier Ethernet Network
  - E-Line: point-to-point service
  - E-LAN: multipoint-to-multipoint service
  - E-Tree: rooted multipoint service
- MEF 51 OVC services: in an operator Carrier Ethernet Network
  - E-Access: point-to-point and multipoint service with UNIs and ENNIs
  - E-Transit: point-to-point and multi-point service with ENNIs

E-Line

E-Line Service Types include Virtual Private Lines (EVPL), Ethernet Private Lines (EPL), Ethernet Internet Access, and TDM Over Ethernet.

EVPL replaces Frame Relay or ATM L2 VPN services. It delivers higher bandwidth, end-to-end services. EVPL enables multiple services (EVCs) to be delivered over a single physical connection (user network interface) to customer premises.

This figure shows an E-Line service type Point-to-Point EVC among the same service provider.

Figure 7-23
E-Line Service type Point-to-Point EVC

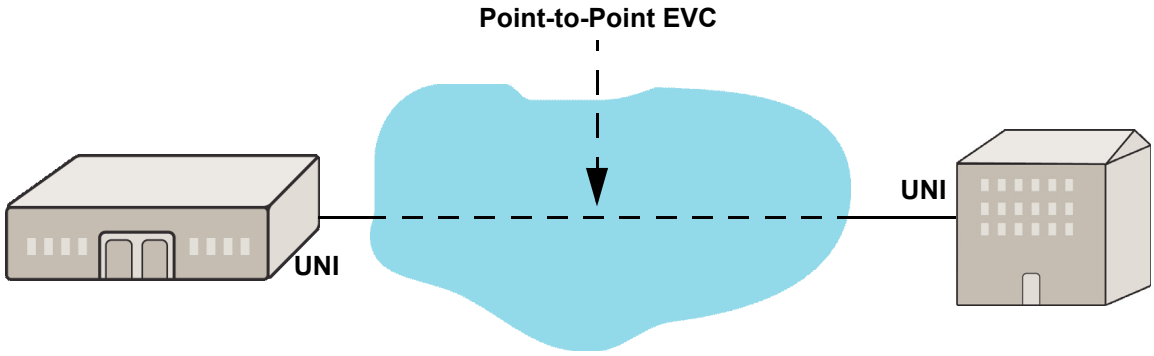

For EPL, the point-to-point EVC provides connectivity between two user network interfaces (UNIs) where all frames sent to the EVC from one UNI are received by the other UNI. EPL provides a port-based service with single

Copyright© 2022 Ciena® Corporation

service EVC across dedicated UNIs to provide site-to-site connectivity. Ethernet over SDH/OTN/MPLS is typically delivered with one customer/ service per user network interface (UNI) port.

E-LAN

E-LAN Service Types include multipoint L2 VPNs, Transparent LAN Service, and Multicast networks. E-LAN Service Types are the foundation for IPTV and Multicast networks. EP-LANs and EVP-LANs use this service type.

For EP-LANs, bridged LAN service is provided across a geographically dispersed network. Each UNI is dedicated to the EP-LAN service. For EVP-LANs, service multiplexing is allowed at each UNI. Multiple LAN services for different customers are provided per each UNI port. An example of an E-LAN Service Type is Internet access and corporate VPN access using one UNI.

This figure shows an E-LAN service type Multi-point-to-Multi-point EVC.

Figure 7-24
E-LAN Service Type for multi-point-to-multi-point EVC

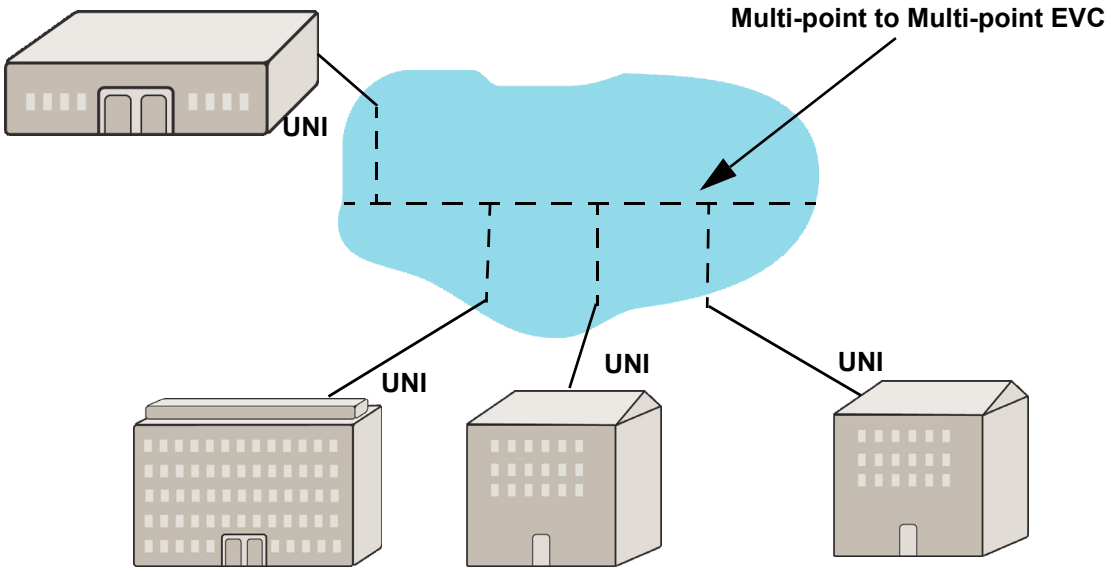

Copyright© 2022 Ciena® Corporation

E-Tree

E-Tree Service Types include rooted multi-point L2 VPNs, broadcast networks, and telemetry networks. The E-Tree Service Type includes EP-Tree and EVP-Tree. Both allow root-to-root and root-to-leaf communication, but not leaf-to-leaf communication. EP-Tree requires dedicated UNIs to the single EP-Tree service. The EVP-Tree lets each UNI support multiple, simultaneous services in a more complex configuration than EP-Tree.

The E-Tree Service Type is used for broadcast applications. For example, the Root UNI can broadcast TV channels to all Leaf UNIs, and the Leaf UNIs can send channel-changing information from each subscriber's box back to the Root UNI.

This figure shows an E-Tree service type rooted multi-point EVC.

Figure 7-25
Rooted multi-point EVC

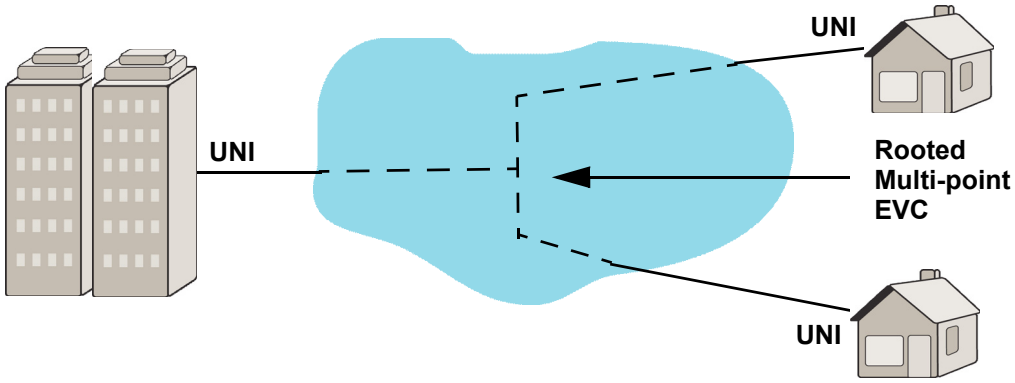

E-Access

E-Access Service Type includes wholesale access services, Access EPL, and Access EVPL. Access providers sell their services to retail. Service providers deliver an EVC-based service to their subscribers. Access EPL and Access EVPL use point-to-point Operator Virtual Connection (OVC) with one UNI endpoint and one ENNI endpoint. An OVC with a UNI and ENNI endpoint is used for Ethernet access services where an Ethernet access provider provides a UNI at the subscriber's premises. In this case, the Ethernet service provider that does not have network facilities to reach the subscriber's site. The site can only be reached through an Ethernet access provider with network facilities that reach the site.

This figure shows an E-Access service type point-to-point EVC among different service providers.

Copyright© 2022 Ciena® Corporation

Figure 7-26
E-Access Service Type for point-to-point EVC

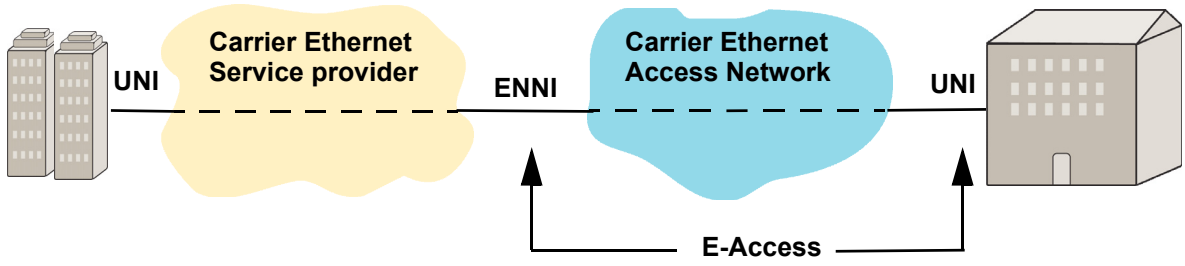

MPLS encapsulation of MEF8 TDM PWs

MPLS encapsulation of MEF8 TDM PWs can be configured by using the MPLS virtual switch component. A VPLS or VPWS virtual switch is used to configure MPLS encapsulation. When the MPLS VS is created and the TDM virtual circuit is created, then the TDM virtual circuit can be added to the MPLS VS with a specified VLAN. The TDM virtual-circuit identifies the physical port and the VLAN identifies which packets to accept into the VS. Adding the TDM virtual circuit creates an MPLS attachment circuit on the physical port.

For double tagging, the CVID is specified when the tdm vc is created. The SVID is specified when the tdm-vc is added to the MPLS VS.

The PWE module and enhanced module TDM ports and TDM SFP ports are supported.

See the procedure"Configuring MPLS encapsulation of MEF8 TDM PWs" on page 7-158.

Copyright© 2022 Ciena® Corporation

Procedure 7-1
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	Create an Ethernet virtual circuit using the provider VLAN. `virtual-circuit ethernet create {vc <vc>} {vlan <Vlan>} [l2-transform <vc-transform-name>]`
4	Create the virtual switch and associate it with the virtual circuit created in step 3. `virtual-switch ethernet create vs <vs> vc <VirtualCircuitEthernetName> encap-fixed-dot1dpri <NUMBER: 0-7>`
5	Disable the LLDP on the UNI port. `lldp set port <port number> disable`
6	Add UNI PORT 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>] [pcp <NUMBER: 0-7>]`

—end—

Copyright© 2022 Ciena® Corporation

Example

This figure shows an example of an EPL provider bridge.

Figure 7-27
Ethernet Private Line

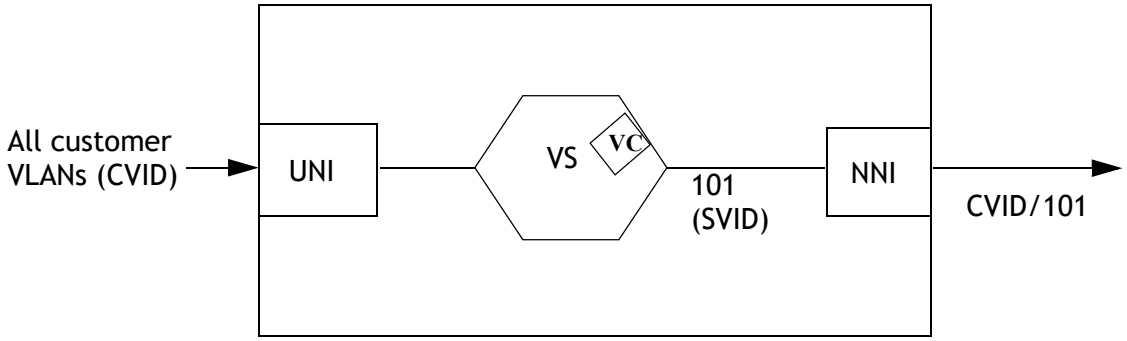

This example configures the EPL provider bridge.

```
> vlan create vlan 101
> vlan add vlan 101 port 10
> virtual-circuit ethernet create vc EVC101 vlan 101
> virtual-switch ethernet create vs EVS100 vc EVC101 encap-fixed-dot1dpri 3
> lldp set port 1 disable
> virtual-switch ethernet add vs EVS100 port 1
```


Copyright© 2022 Ciena® Corporation

Procedure 7-2
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 <vc>} {vlan <Vlan>} [l2-transform <vc-transform-name>]`
4	Create Ethernet virtual switches using the Ethernet virtual circuit. `virtual-switch ethernet create vs <vs> vc <VirtualCircuitName>`
5	Disable the LLDP on the UNI port. `lldp set port <port number> disable`
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>] [pcp <NUMBER: 0-7>]`

—end—

Example

This figure shows an example of an EVPL provider bridge.

Copyright© 2022 Ciena® Corporation

Figure 7-28
EVPL Provider Bridging

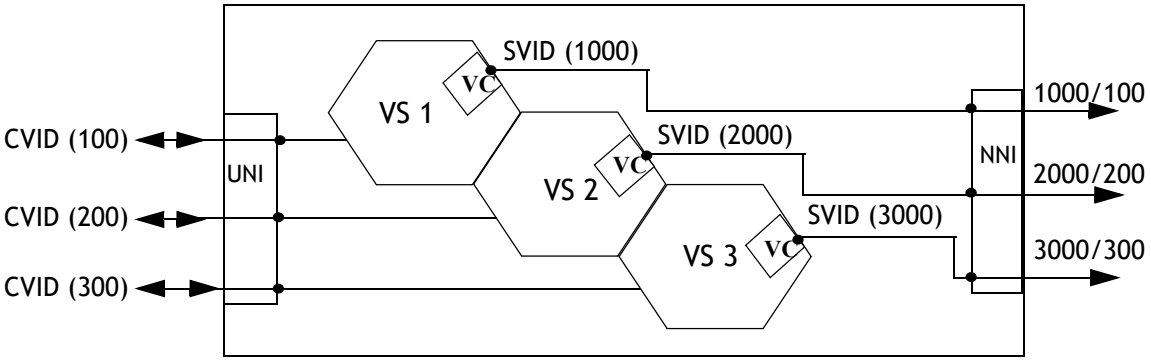

This example creates the configuration.

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
```

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
```

Disable the LLDP on the UNI port.

```
> lldp set port 1 disable
```

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© 2022 Ciena® Corporation

Procedure 7-3
Configuring a virtual switch to emulate a VLAN on a UNI port

Configure a virtual switch to emulate a VLAN on a UNI port.

Step	Action
1	Set the virtual switch transform mode for all UNI ports to i-stamp:push,e-match-pop:stamp: `port set port <PortNameList> vs l2-transform i-stamp:push,e-match-pop-stamp` where port <PortNameList> — is the name of the ports.
2	If required, reserve VLANs for the virtual switch: `virtual-switch add reserved-vlan <VlanIDList>` where reserved-vlan <VlanIDList> — is the name of the VLANs.
3	Create the VLAN for the traffic to be switched: `vlan create vlan <VlanIDList>` where vlan <VlanIDList> — is the name of the VLANs.
4	Disable the LLDP on the UNI port. `lldp set port <port number> disable` where port number — is the port number.
5	Add the NNI port(s): `vlan add vlan <VlanIDList> port <PortNameList>` where vlan <VlanIDList> — is the name of the VLANs. port <PortNameList> — is the name of the port.

Copyright© 2022 Ciena® Corporation

6 Create the virtual circuit for the virtual switch using the traffic's VLAN:

```
virtual-circuit ethernet create vc <vc> vlan <VlanIDList>
```

where

vc <vc>	is the name of the Ethernet Virtual Circuit.
vlan <VlanIDList>	is the name of the VLANs.

7 Create the virtual switch using the new virtual circuit:

```
virtual-switch ethernet create vs <vs> vc
<VirtualCircuitEthernetName>
```

where

vs <vs>	is the name of the Virtual Switch.
vc <VirtualCircuitEthernet Name>	is the name of the Ethernet Virtual Circuit.

8 Create an EVPL virtual switch member with a translate tag of zero on the UNI port, with the vlan attribute equal to the VLAN to emulate:

```
virtual-switch ethernet add vs <VirtualSwitchName> port
<PortNameList> vlan <VlanIDList> translate-tag 0
```

where

vs <VirtualSwitchName>	is the name of the Virtual Switch.
port <PortNameList>	is the name of the UNI port.
vlan <VlanIDList>	is the name of the VLAN to emulate.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example shows how a virtual switch emulates a VLAN on a UNI port.

Figure 7-29
VLAN emulation on a UNI port example

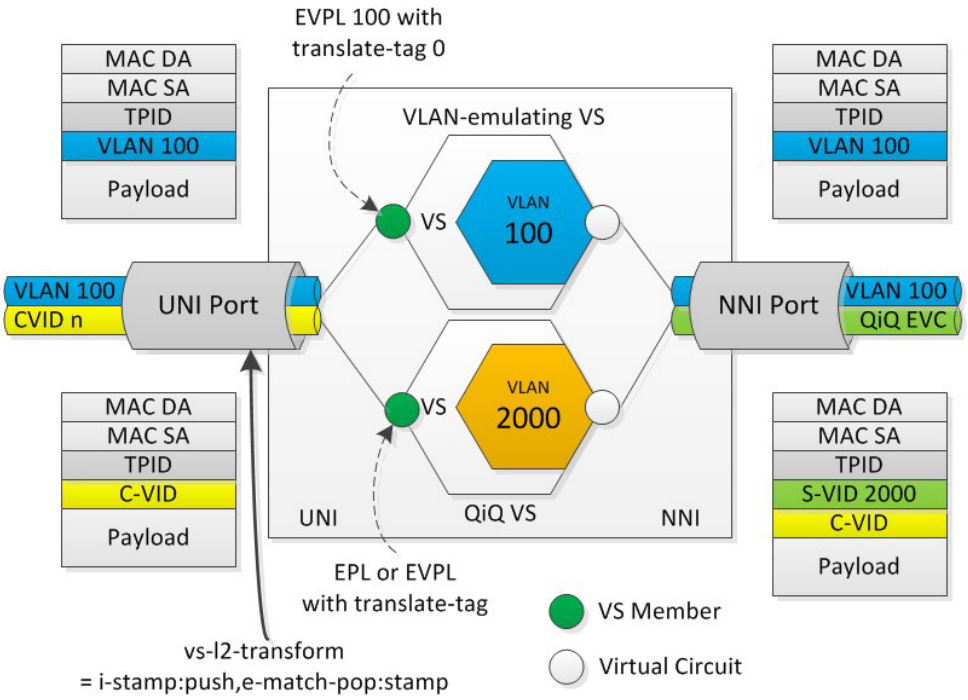

This example shows how to configure two virtual switches to emulate a VLAN on a UNI port. The UNI port is 1, and the NNI port is 24.

```
port set port 1 vs-l2-transform i-stamp:push,e-match-pop:stamp
vlan create vlan 100

vlan add vlan 100 port 24

lldp set port 1 disable

virtual-circuit ethernet create vc vc100 vlan 100
virtual-switch ethernet create vs vs100 vc vc100
virtual-switch ethernet add vs vs100 port 1 vlan 100 translate-tag 0
vlan create vlan 2000 port 24
virtual-switch ethernet create vs vs2000 vc vc2000
virtual-switch ethernet add vs vs2000
port 1 vlan 2000 translate-tag <CVID>
```


Copyright© 2022 Ciena® Corporation

Procedure 7-4
Configuring the fixed encapsulation priority value

Configure the fixed encapsulation priority value. You can configure the fixed encapsulation priority value for

- a virtual switch on creation
- an existing virtual switch

Step	Action

To configure the fixed encapsulation priority value for a virtual switch on creation

1 Configure the fixed encapsulation priority value for a virtual switch on creation:

```
virtual-switch ethernet create vs <vs> encap-fixed-
dot1dpri <NUMBER:0-7>
```

where

vs <vs>	is the virtual switch name.
encap-fixed-dot1pri <NUMBER: 0-7>	is the encapsulation fixed dot1dpri.

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-interface
<interface-name>]
```

where

vs <vs>	is the virtual switch name.
{decap-cos-policy <fixed | leave | tunnel-inherit | vc-inherit}	is the decapsulation CoS policy. The leave policy is not supported for PBT in native mode.
decap-fixed-dot1dpri <NUMBER: 0-7>	is the decapsulation fixed dot1dpri.
description <String[128]>	sets the VS description.

Copyright© 2022 Ciena® Corporation

where

encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit>	is the encapsulation CoS policy.
encap-fixed-dot1pri <NUMBER: 0-7>	is the encapsulation fixed dot1dpri.
subscriber-dot1dpri-policy <leave | provider-inherit>	is the subscriber dot1d marking behavior.
vc <Virtual Circuit Name>	is the Ethernet virtual circuit.
ip-interface <interface-name>	is the IP interface name.

Copyright© 2022 Ciena® Corporation

Procedure 7-5
Configuring UNI member statistics

Configure the UNI member statistics to obtain these statistics from the device:

- Rx statistics
- Accepted Rx statistics
- Tx statistics
- Rx throughput
- Accepted Rx throughput
- Tx throughput

You cannot set untagged statistics collection for a UNI member of type EVPL if another

- EVPL with the same port is already present in any virtual switch where untagged statistics collection is enabled.
- EPL with the same port is already present in any virtual switch where untagged statistics collection is enabled with statistics collection on.

You cannot set untagged statistics collection for a UNI member of type EPL if another EVPL with the same port is already present in any virtual switch where the untagged statistics collection is enabled with statistics collection on.

Step	Action
1	Create an SVLAN: `vlan create vlan <vlan>` where vlan <vlan> is SVLAN name.
2	Add the NNI port in the SVLAN: `vlan add vlan 100 port <port>` where port <port> is the port number.
3	Create a virtual circuit: `virtual-circuit ethernet create vc <vc> vlan <vlan>` where vc <vc> is the virtual circuit name.

Copyright© 2022 Ciena® Corporation

4 Create a virtual switch:

```
virtual-switch ethernet create vs <vs> vc <vc>
```

where

vs <vs>	is the virtual circuit name.

5 Add the port as an EVPL member in the virtual switch:

```
virtual-switch ethernet add vs <vs> port <port> vlan
<vlan>
```

6 Add the port as an EPL member in the virtual switch:

```
virtual-switch ethernet add vs <vs> port <port>
```

7 Configure a traffic profile for an EVPL member of a virtual switch:

```
traffic-profiling set port <port> mode advanced
traffic-profiling set port <port> classifier-mode narrow
traffic-profiling standard-profile create port <port>
profile <NUMBER: 1..4096> name <String[1..15]> cir <Kbps>
pir <Kbps> cbs <KBYTES: 0..262144> ebs <kbytes:
0..262144> vlan <number: 1..4095>
traffic-profiling enable port <port>
traffic-profiling enable
```

where

port <port>	is the port name.
profile <NUMBER: 1..4096>	is the traffic profile number.
name <String[1..15]>	is the traffic profile name.
cir <Kbps>	is the committed information rate in Kbps of this traffic meter, rounded to fit hardware.
pir <Kbps>	is the peak information rate in Kbps of this traffic meter, rounded to fit hardware.
cbs <KBYTES: 0..262144>	is the meter committed burst size in Kbytes, rounded to the nearest size supported in hardware.
ebs <kbytes: 0..262144>	is the meter excess burst size in Kbytes, rounded to the nearest size supported in hardware.

Copyright© 2022 Ciena® Corporation

8 Configure a traffic profile for an EPL member of a virtual switch:

```
traffic-profiling set port <port> mode advanced
traffic-profiling set port <port> classifier-mode narrow
traffic-profiling standard-profile create port <port>
profile <profile> name <String[1..15]> cir <Kbps> pir
<Kbps> cbs <KBYTES: 0..262144> ebs <kbytes: 0..262144>
default-profile-true
traffic-profiling enable port <port>
traffic-profiling enable
```

9 Configure a traffic profile for an EVPL member of a virtual switch with untagged statistics:

```
traffic-profiling set port <port> mode advanced
traffic-profiling set port <port> classifier-mode narrow
traffic-profiling standard-profile create port <port>
profile <profile> name <String[1..15]> cir <Kbps> pir
<Kbps> cbs <KBYTES: 0..262144> ebs <kbytes: 0..262144>
vlan <1..4095> untagged
traffic-profiling enable port <port>
traffic-profiling enable
```

where

vlan <number: 1..4095>	is the VLAN-ID to monitor.

10 Enable statistics collection for each UNI member:

```
virtual-switch ethernet set port <port> vs <vs> vlan
<vlan> statistics on include-untagged-statistics on rx-
statistics-profile <TrafficProfileName>
```

where

vs <vs>	is the virtual switch name.
statistics-profile <TrafficProfileName>	is the traffic profile that you want to monitor.

—end—

Example

In the UNI member statistics example figure, port 1 is added in the virtual switch as EVPL with C-TAG 10 and 11. Port 2 is added in virtual switch as EPL member. Traffic-profile needs to be configured for each EVPL and EPL member. Statistics collection is enabled for each EVPL/EPL member. Untagged traffic is transmitted out of port 1 and is counted as part of the transmitted statistics of an EVPL member with C-TAG 10.

Copyright© 2022 Ciena® Corporation

Figure 7-30
UNI member statistics example

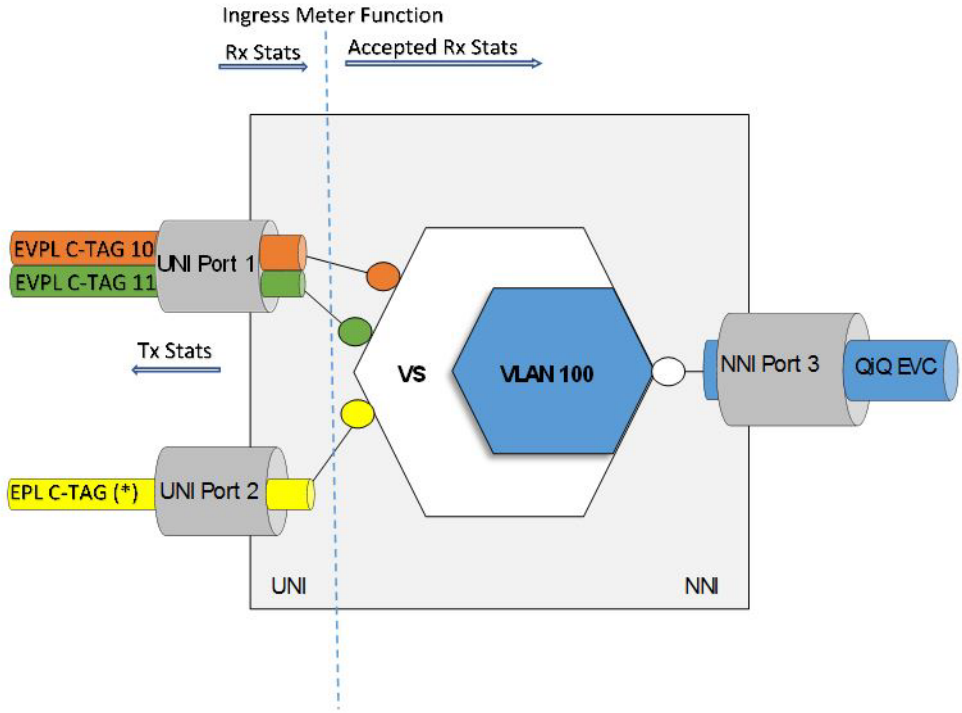

This example configures UNI member statistics.

```
vlan create vlan 100
vlan add vlan 100 port 3
virtual-circuit ethernet create vc vc1 vlan 100
virtual-switch ethernet create vs vs1 vc vc11
virtual-switch ethernet add vs vs1 port 1 vlan 10
virtual-switch ethernet add vs vs1 port 1 vlan 11
virtual-switch ethernet add vs vs1 port 2
traffic-profiling set port 1 mode advanced
traffic-profiling set port 2 classifier-mode narrow
traffic-profiling standard-profile create port 1 profile 1 name TpStatP1V10
cir 1000 pir 1000 cbs 1000 ebs 1000 vlan 10 untagged
traffic-profiling standard-profile create port 1 profile 2 name TpStatP1V11
cir 2000 pir 2000 cbs 2000 ebs 2000 vlan 11
traffic-profiling enable port 1
traffic-profiling set port 2 mode advanced
traffic-profiling set port 2 classifier-mode narrow
traffic-profiling standard-profile create port 2 profile 3 name TpStatP2 cir
3000 pir 3000 cbs 3000 ebs 300 default-profile true
traffic-profiling enable port 2
traffic-profiling enable
virtual-switch ethernet set port 1 vs vs1 vlan 10 statistics on include-
untagged-statistics on rx-statistics-profile TpStatP1V10
virtual-switch ethernet set port 1 vs vs1 vlan 11 statistics on
rx -statistics-profile TpStatP1V11
virtual-switch ethernet set port 2 vs vs1 statistics on rx-statistics-profile
TpStatP2 include-untagged-statistics on
```


Copyright© 2022 Ciena® Corporation

Procedure 7-6
Clearing UNI member statistics

You can clear

- UNI member statistics of a single UNI member
- UNI member statistics of all the UNI members of a virtual switch

Step	Action

To clear UNI member statistics of a single UNI member

1 Clear UNI member statistics of a single UNI member.

```
virtual-switch ethernet clear port <port> vs <vs> vlan
<vlan> {statistics}
```

```
virtual-switch ethernet clear port <port> vs <vs>
{statistics}
```

where

port <port>	is the subscriber port.
vs <vs>	is the virtual switch name.
vlan <vlan>	is the VLAN ID list.
statistics	is the subscriber member statistics.

To clear UNI member statistics of all the UNI members of a virtual switch

2 Clear UNI member statistics for UNI members of a virtual switch.

```
virtual-switch ethernet clear vs <vs> {statistics}
```

—end—

Example

This example clears UNI member statistics of a single UNI member.

```
virtual-switch ethernet clear port 1 vs vs1 vlan 10
statistics
```

```
virtual-switch ethernet clear port 2 vs vs1 statistics
```

This example clears UNI member statistics of all the UNI members of a virtual switch.

```
virtual-switch ethernet clear vs vs1 statistics
```


Copyright© 2022 Ciena® Corporation

Procedure 7-7
Displaying UNI member information

Display UNI member

- statistics
- throughput

Step	Action

To display UNI member statistics

1 Display UNI member statistics:

```
virtual-switch ethernet show port <port> vs <vs> vlan
<vlan> statistics
```

where

port <port>	is the port name.
vs <vs>	is the virtual switch name.
vlan <vlan>	is SVLAN name.

To display UNI member throughput

2 Display UNI member throughput:

```
virtual-switch ethernet show port <port> vs <vs> vlan
<vlan> throughput scale kilo
```

—end—

Copyright© 2022 Ciena® Corporation

Example

This output shows statistics for the virtual-switch ethernet show command.

```
virtual-switch ethernet show port 1 vs vs1 vlan 10 statistics

+--------- VIRTUAL SWITCH MEMBER STATISTICS ---------+
| Port   | Subscriber | Statistics   |        |         |
|        | VLAN       | Type         | Bytes  | Packets |
+--------+------------+--------------+--------+---------+
| 1      | 10         | Rx           | 0      | 0       |
|        |            | Rx Accepted  | 0      | 0       |
|        |            | Tx           | 0      | 0       |
+--------+------------+--------------+--------+---------+
```

This output is displayed throughput for the virtual-switch ethernet show command.

```
virtual-switch ethernet show port 1 vs vs1 vlan 10 throughput scale kilo

+------------------------- VIRTUAL SWITCH MEMBER THROUGHPUT --------------+
| Statistics         | Current Value | Delta Value     | Rate (kbps & kpps) |
+--------------------+---------------+-----------------+--------------------+
| Time               |    7:05:25:22 |  0:00:00:00.73  |                    |
+--------------------+---------------+-----------------+--------------------+
| Rx Bytes           |             0 |               0 |                  0 |
| Accepted Rx Bytes  |             0 |               0 |                  0 |
| Tx Bytes           |             0 |               0 |                  0 |
| Rx Pkts            |             0 |               0 |                  0 |
| Accepted Rx Pkts   |             0 |               0 |                  0 |
| Tx Pkts            |             0 |               0 |                  0 |
+--------------------+---------------+-----------------+--------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 7-8
Determining egress resources to collect Tx statistics

Sub-port statistics and UNI member statistics share resources to collect Tx statistics. This procedure lets you determine the resource usage.

Step	Action
1	Determine resource usage: `resource-manager show feature egress-tx-statistics`

—end—

Example

This example shows the output from the resource-manager show feature egress-tx statistics command.

```
> resource-manager show feature egress-tx-statistics

+---------- EGRESS-TX-STATISTICS RESOURCE INFO ----------+
|              | Classifiers |   Meters    |  Counters   |
+--------------+-------------+-------------+-------------+
| Available    |         256 |         256 |         256 |
| Used         |          20 |           0 |          20 |
| Free         |         236 |         256 |         236 |
+--------------+-------------+-------------+-------------+
+--------- EGRESS-TX-STATISTICS RESOURCE USAGE ----------+
|              | Classifiers |   Meters    |  Counters   |
+--------------+-------------+-------------+-------------+
| Sub-Port     |           5 |           0 |           5 |
| VS Member    |          15 |           0 |          15 |
+--------------+-------------+-------------+-------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 7-9
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>] [pcp <NUMBER: 0-7>]` where

vs <vs>	is the virtual switch name.
port <port>	is the subscriber port.
vlan <VLAN list>	is the subscriber VLAN ID.
ip-interface <interface-name>	is the interface name.
encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit | port inherit | vs-inherit>>	is the encapsulation CoS policy.
encap-fixed-dot1pri <NUMBER: 0-7>	is the encapsulation fixed dot1dpri.
statistics <on|off>	is the statistics control for Virtual Switch Member.
translate-tag <NUMBER: 0..4094>	is the subscriber VLAN translate stamp.
[pcp <NUMBER: 0-7>]	is the egress CVID priority (pcp) value on packets egressing the vs-member IP interface. The default value is 0.

Copyright© 2022 Ciena® Corporation

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-interface
<interface-name>]
```

where

vs <vs>	is the virtual switch name.
{decap-cos-policy <fixed | leave | tunnel-inherit | vc-inherit}	is the decapsulation CoS policy. The leave policy is not supported for PBT in native mode.
decap-fixed-dot1dpri <NUMBER: 0-7>	is the decapsulation fixed dot1dpri.
description <String[128]>	sets the VS description.
encap-cos-policy <dot1dpri-inherit | fixed | ip-prec-inherit | phbg-inherit>	is the encapsulation CoS policy.
encap-fixed-dot1pri <NUMBER: 0-7>	is the encapsulation fixed dot1dpri.
subscriber-dot1dpri-policy <leave | provider-inherit>	is the subscriber dot1d marking behavior.
vc <Virtual Circuit Name>	is the Ethernet virtual circuit.
ip-interface <interface-name>	is the IP interface name.

—end—

Example

This 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
```

This example sets the CoS policy for existing VS members.

```
virtual-switch ethernet set port 5 vs vs150 encap-cos-policy vs-
inherit
```


Copyright© 2022 Ciena® Corporation

Procedure 7-10
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:*** The untagged-data-vid attribute does not apply to MPLS virtual switches.

Step	Action
1	Create a provider VLAN (SVID): `vlan create vlan <VlanIdList> [name <String[31]>]`
2	Create an Ethernet virtual circuit using the provider VLAN. `virtual-circuit ethernet create vc <vc> {vlan <Vlan>}[l2-transform <vc-transform-name>]`
3	Create the virtual switch and associate it with the virtual circuit: `virtual-switch ethernet create vs <vs> vc <vc>`
4	Disable the LLDP on the UNI port. `lldp set port <port number> disable`
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>] [pcp <NUMBER: 0-7>]`

Copyright© 2022 Ciena® Corporation

6 Set UNI port attributes:

```
port set port <PortNameList> acceptable-frame-type <all |
tagged-only | untagged-only> untagged-data-vs
<VirtualSwitchName> untagged-data-vid <VlanId> fixed-
rcos <NUMBER: 0-7]
```

where

port <PortNameList>	is the port list
acceptable-frame-type <all | tagged-only | untagged-only>	sets the acceptable frame types mode as follows: • all - all frames are allowed to ingress, regardless of their tag status. • tagged-only - only tagged frames are allowed to ingress. • untagged-only - only untagged frames are allowed to ingress. When acceptable frame types are configured on a non-virtual switch UNI, untagged0only mode filters out all three common TPIDs of 0x8100, 0x8898, and 0x9100. For a virtual switch UNI, only the TPID of 0x8100 is recognized as a tagged frame.
untagged-data-vs <VirtualSwitchName>	is the virtual switch for untagged data frames.
untagged-data-vid <VlanId>	is the VLAN ID for untagged data frames.
fixed-rcos <NUMBER: 0-7]	is the Fixed Resolved CoS value.

—end—

Example

This example sets these 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-circuit ethernet create vc vc555 vlan 555
> virtual-switch ethernet create vs vsX vc vc555
> lldp set port disable port 1
> 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© 2022 Ciena® Corporation

Procedure 7-11
Setting the virtual switch L2 transform action on a port

Set the virtual switch L2 transform action on a port.

To enable transform actions other than the default, assign resources for the feature.

> ***Note:*** For more information regarding resource management, refer to "Hardware resource management" on page 4-1.

Step	Action
1	Set the VS 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 counter resources to 512. `resource-manager pool set resource counter feature vs-enhanced-l2-transform count 512`
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© 2022 Ciena® Corporation

Procedure 7-12
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 <vc>} {vlan <Vlan>}`
5	Create a virtual switch instance and associate a virtual circuit: `virtual-switch ethernet create vs <vs> [vc <vc>]`
6	Disable the LLDP on the UNI port. `lldp set port <port number> disable`
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>] [pcp <NUMBER: 0-7>]`
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. `vlan add vlan <VlanList> port <PortNameList>`

Copyright© 2022 Ciena® Corporation

11 Create a static Ethernet virtual circuit for the specified VLAN object:

```
virtual-circuit ethernet create {vc <vc>} {vlan <Vlan>}
```

12 Create a virtual switch instance and associate a virtual circuit:

```
virtual-switch ethernet create vs <vs> [vc <vc>]
```

13 Disable the LLDP on the UNI port.

```
lldp port <port number> disable
```

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>] [pcp <NUMBER: 0-7>]
```

—end—

Copyright© 2022 Ciena® Corporation

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 virtual switch L2 transform action on a port".

Figure 7-31
VS L2 transform: i-push,e-pop:stamp for Q-in-Q

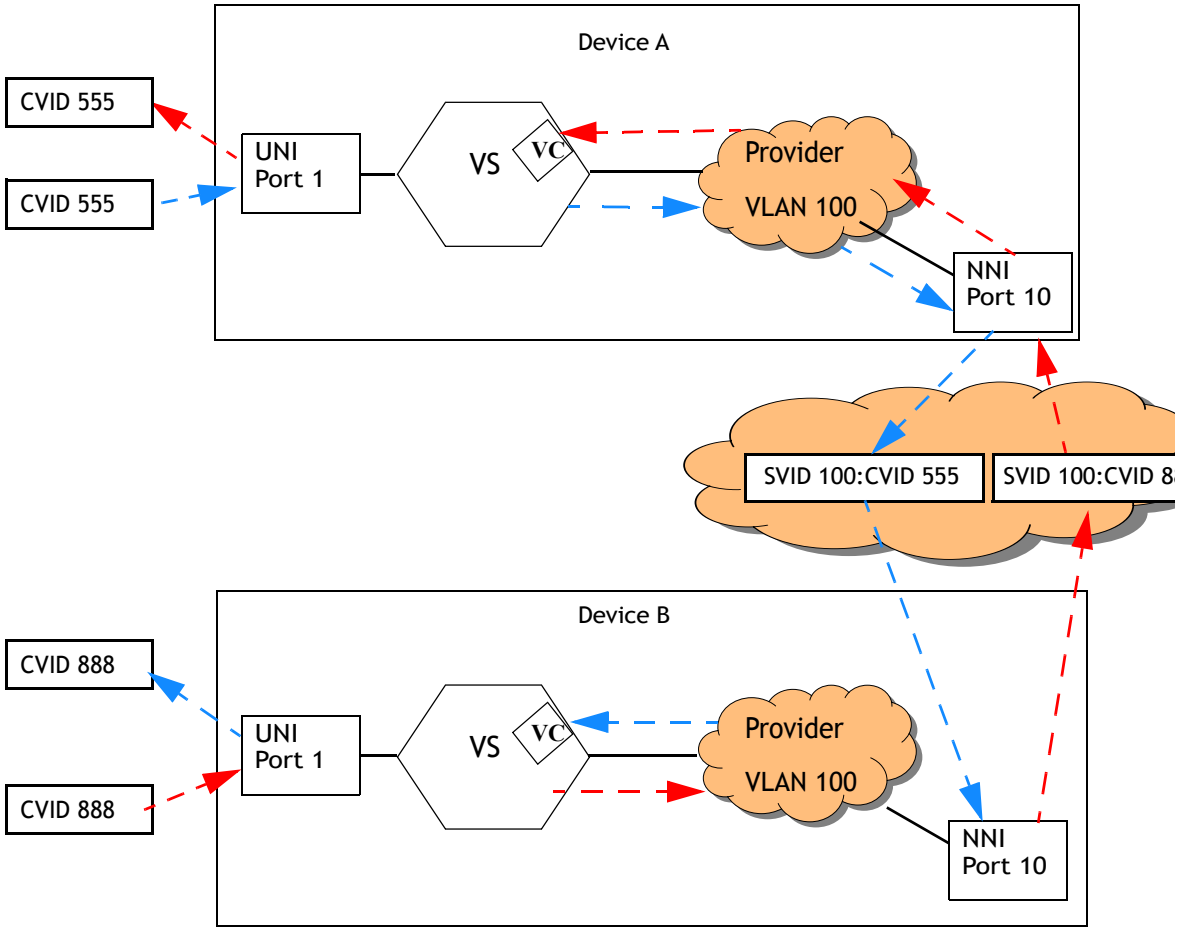

Copyright© 2022 Ciena® Corporation

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 ethernet create vs vs100 vc vc1
lldp set port port 1 disable
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 ethernet create vs vs100 vc vc1
lldp set port port 1 disable
virtual-switch ethernet add vs vs100 port 1 vlan 888
```


Copyright© 2022 Ciena® Corporation

Procedure 7-13
Creating an i-stamp:push,e-match-pop:stamp configuration

Create an i-stamp:push,e-match-pop:stamp configuration. Untagged data virtual switch and control virtual switch are supported.

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 <vc>} {vlan <Vlan>}[l2-transform <VC transform name>]`
5	Create a virtual switch instance and associate a virtual circuit: `virtual-switch ethernet create vs <vs> [vc <vc>]`
6	Disable the LLDP on the UNI port. `lldp set port <port number> disable`
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>] [pcp <NUMBER: 0-7>]`

—end—

Example

These examples show

- “When translate-tag is non-zero”
- “When translate-tag is zero”

Copyright© 2022 Ciena® Corporation

When translate-tag is non-zero

In this scenario, frames with CVID 555, 600, or 777 classify to the UNI port 1. On 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. For 3903, 3903x, 3904, 3905,3924, 3906, 3928 platforms, this configuration requires the allocation of resources as described in "Setting the virtual switch L2 transform action on a port".

Figure 7-32
VS L2 transform: i-push,e-pop for Q-in-Q

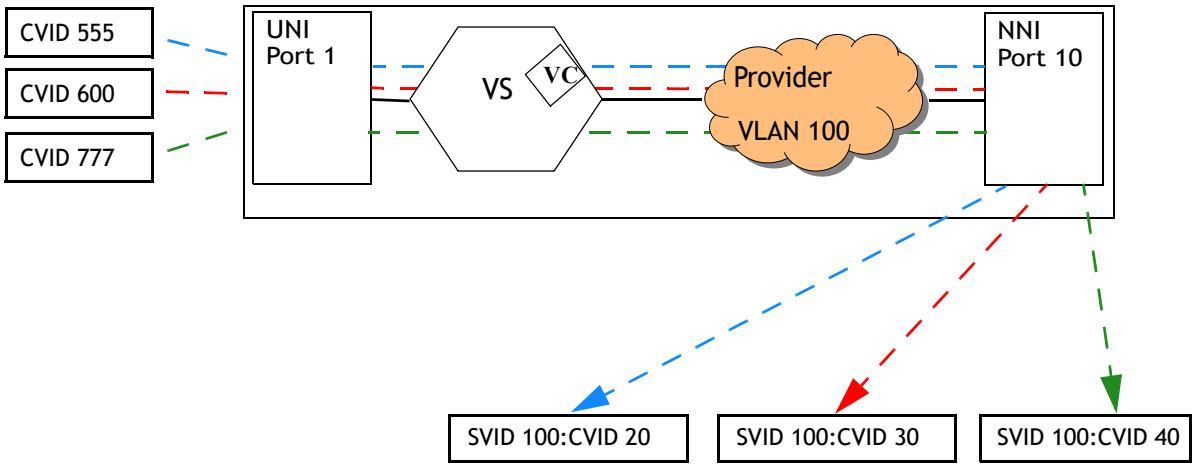

This example creates the configuration.

```
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 ethernet create vs vs100 vc vc100
lldp set port 1 disable
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
```


Copyright© 2022 Ciena® Corporation

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, and the virtual switch has a reserved VLAN, 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.

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 virtual switch L2 transform action on a port”.

This 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 ethernet create vs X vc X
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© 2022 Ciena® Corporation

Procedure 7-14
Configuring flood traffic from a NNI port toward a UNI port

You can limit traffic from a NNI port toward a UNI port in an i:stamp:push,e-match-pop:stamp configuration. The vs-egress-filter attribute is used to limit flooding traffic coming in from a NNI port toward a UNI port.

Step	Action
1	Create an SVLAN. `vlan create vlan <String[31]>` where vlan <String[31]> is the VLAN ID.
2	Set the UNI port to the enhanced transform mode. `port set port <PortNameList> vs-l2-transform <i-push,epop | i-push,e-pop;stamp | i-stamp:push,e-match-pop:stamp>}` where port <PortNameList> — the ports to set. vs-l2-transform <i-push,e-pop | i-push,e-pop:stamp | i-stamp:push,e-match-pop:stamp>} — sets the L2 transform action for virtual switch subscribers. For this procedure, i-stamp:push,e-match-pop:stamp must be used.
3	Add the SVLAN to the NNI port. `vlan add vlan <String[31]> port <PortNameList>` where vlan <String[31]> — the VLAN being added to the NNI port. port <PortNameList> — the ports to add.

Copyright© 2022 Ciena® Corporation

4 Set the UNI port attributes with vs-egress-filter enabled.

```
port set port <PortList> acceptable-frame-type
<all|tagged-only|untagged-only> untagged data-vid
<virtual switch name> vs-egress-filter <on|off>
```

where

port <PortNameList>	the ports to set.
acceptable-frame-type <all|tagged-only|untagged-only>	sets the acceptable frame type. For this procedure, untagged-only must be used.
untagged data-vid <virtual switch name>	sets the virtual switch for untagged data frames.
vs-egress-filter <on|off>	enables or disables the vs-egress-filter.

5 Configure the virtual circuit using SVLAN.

```
virtual-circuit ethernet create vc <vc> vlan <vlan name>
```

where

vc <vc>	is the virtual circuit name.
vlan <vlan name>	is the provider VLAN.

6 Configure the virtual switch with a virtual circuit.

```
virtual-switch ethernet create vs <vs> vc <vc>
```

where

vs <vs>	the virtual switch name.
vc <vc>	is the virtual circuit name.

7 Add UNI ports as VS EPL members.

```
virtual-switch ethernet add vs <vs> port <PortNameList>
```

where

vs <vs>	is the virtual switch name.
port <PortNameList>	are the UNI port numbers.

Copyright© 2022 Ciena® Corporation

8 Set the UNI port untagged data to VS.

```
port set port <NUMBER> untagged-data-vs <virtual switch
name
```

where

port <port number>	is the UNI port number.
untagged-data-vs <virtual switch name>	sets the virtual switch for untagged data frames.

—end—

Example

This example configures flood traffic from a NNI port to a UNI port.

```
vlan create vlan 1000
port set port 1 vs-l2-transform i-stamp:push,e-match-pop:stamp
port set port 2 vs-l2-transform i-stamp:push,e-match-pop:stamp
vlan add vlan 1000 port 4
port set port 1 acceptable-frame-type untagged-only untagged-data-vid 10 vs-
egress-filter on
port set port 2 acceptable-frame-type untagged-only untagged-data-vid 20 vs-
egress-filter on
virtual-circuit ethernet create vc vc1000 vlan 1000
virtual-switch ethernet create vs vs1000 vc vc1000
virtual-switch ethernet add vs vs1000 port 1
virtual-switch ethernet add vs vs1000 port 2
port set port 1 untagged-data-vs vs1000
port set port 2 untagged-data-vs vs1000
```


Copyright© 2022 Ciena® Corporation

Procedure 7-15
Configuring a virtual circuit transform

Configure a virtual circuit transform according to the network plan. A VC is where a configured instance of a VC transform is assigned to one or more VCs. A VC transform provides a VC transform service to a VC that is assigned to the profile, and all associated entities of the VC.

Step	Action
1	Configure a VC transform: `virtual-circuit transform set l2-transform <VC transform name> {dei <NUMBER: 0-1>}{pcp <NUMBER: 0-7} {tpid <8100|9100|88A8>} {vid <NUMBER: 1-4094>}` **—end—**

Example

This example configures a VC transform.

```
virtual-circuit transform set l2-transform ptag1000 dei1 pcp 5 tpid 9100 vid
4094
```


Copyright© 2022 Ciena® Corporation

Procedure 7-16
Configuring VC transform support on the 3928 and 5160 switches

Configure VC transform support on the 3928 and 5160 switches before configuring a virtual circuit transform service on the 3928 and 5160 switches.

By default, the 3928 and 5160 switches do not have enough internal resources to perform VC transform operations. To use VC transform on the 3928 and 5160, one front panel port is configured as a system-reserved port. The system-reserved port can only be used by the VC transform feature.

Before provisioning a port as a system-reserved port, make sure the port is in the default state, that is, it is not configured against any services or features. In addition, none of the port attributes can be modified.

When choosing the port to be configured as the system-reserved port, consider these restrictions:

- port 20 is reserved for OAM services, and cannot be configured as a system-reserved port
- the port cannot be added to a link aggregation group (LAG) and a LAG member cannot be configured as a system-reserved port
- a port used in a TDM port cannot be configured as a system-reserved port, and a system-reserved port cannot be used when creating a TDM port
- an NNI port in a VC transform service cannot be configured as a system-reserved port, and a system-reserved port cannot be used as a NNI port in a VC transform service

After you configure a system-reserved port on the 3928 and 5160 switches, perform "Configuring a virtual circuit transform service" on page 7-97.

Step	Action
1	Make sure the port you have chosen to configure as the system-reserved port is in the default state and none of the port attributes have been modified from the default: `port show port <port>`

Copyright© 2022 Ciena® Corporation

2 Configure the selected port as a system-reserved port:

```
chassis system-rsvd-port add port <port> feature vc-
transform
```

where

port <port>	is the system-reserved port you are adding.
feature vc-transform	is the feature vc-transform that you are configuring the system-reserved port for.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 7-17
Configuring a virtual circuit transform service

A VC transform service for a virtual switch or a virtual switch cross connect is where a VS or VC uses a VC transform service and can push or pop an extra VLAN tag on the NNI port.

This procedure assumes that adequate VC transform resources have been allocated in the resource manager.

Step	Action
1	Create a VLAN to be used as the VLAN ID in the S-Tag of the VC transform service: `vlan create vlan <vlan>` where <vlan> is the VLAN to be used as the VLAN ID in the S-Tag of the VC transform service.
2	Add port members to the created VLAN. These members represent NNI ports on the VC transform service. `vlan add vlan <vlan> port <port>` where <vlan> is the created VLAN. <port> is the port members being added to the created VLAN.
3	Create a VC transform which identifies the outer most P-Tag that is pushed onto frames egressing the NNI port: `virtual-circuit transform create l2-transform <VC transform name> {vid <NUMBER: 1-4094>} [dei <NUMBER: 0-1>] [pcp <NUMBER: 0-7>][tpid <8100|9100|88A8>]` where L2-transform <VC transform name> is the transform name. vid <NUMBER: 1-4094> is the VLAN ID. dei <NUMBER: 0-1> is the drop eligible indication. pcp <NUMBER: 0-7> is the priority code point. tpid <8100_9100|88A8> is the VLAN Ethertype.

Copyright© 2022 Ciena® Corporation

4 Create a virtual circuit using the L2 transform created and the SVLAN. This object represents the service endpoint for this particular switch:

```
virtual-circuit ethernet create vc <vc> vlan <vlan name>
l2-transform <VC-transform name>
```

where

vc <vc>	is the virtual circuit name.
vlan <vlan name>	is the VLAN name.
l2-transform <VC-transform name>	is the transform name.

Associating the configured VC transform service to a virtual switch

5 Create a virtual switch using the virtual circuit created.

```
virtual-switch ethernet create vs <vs> vc <vc>
```

6 Add virtual switch members to the newly created virtual switch:

```
virtual-switch ethernet add vs <vs> port <port> vlan <vlan
ID>
```

Associating the configured VC transform service to a virtual switch cross connect

7 Create a virtual switch cross-connect using the virtual circuit created, and an existing configured TDM virtual circuit. This virtual-switch cross-connect represents the VC transform service.

```
virtual-switch cross-connect create xc <String[31]> tdm-
vc <TDM VC Name> eth-vc <Ethernet VC Name>
```

where

xc <String[31]>	is the cross-connect name.
tdm-vc <TDM VC Name>	is the TDM VC name.
<Ethernet VC Name>	is the VC created in step 4.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example configures a VC transform service.

```
vlan create vlan 100
vlan add vlan 100 port 2
virtual-circuit transform create l2-transform ptag1000 vid 1000 pcp 2
virtual-circuit ethernet create vc vc100 vlan 100 l2-transform ptag1000
```

This example associates the configured VC transform service to a virtual switch:

```
virtual-switch ethernet create vs vs100 vc vc100
virtual-switch ethernet add vs vs100 port 1 vlan 10
virtual-switch ethernet add vs vs100 port 1 vlan 11
```

This example associates the configured VC transform service to a virtual switch cross connect:

```
virtual-switch cross-connect create xc xc100 tdm-vc
```


Copyright© 2022 Ciena® Corporation

Procedure 7-18
Configuring a remote management interface on a VC VLAN on a VC transform

You can configure a remote management interface on a VC VLAN on a VC transform by using a static IP address. this procedure results in management traffic being double tagged on the NNI port (VC port).

Step	Action
1	Disable the DHCP IPv4 client: `dhcp client disable`
2	Disable the DHCP IPv6 client: `dhcpv6 client disable`
3	Create a VLAN to be used as the VLAN ID in the S-Tag of the VC transform service. `vlan create vlan <vlan>` where <vlan> is the VLAN to be used as the VLAN ID in the S-Tag of the VC transform service.
4	Add port members to the created VLAN. These members represent NNI ports on the VC transform service. `vlan add vlan <vlan> port <port>` where <vlan> is the created VLAN. <port> is the port members being added to the created VLAN.
5	Configure the remote management interface on the VLAN. `interface remote set vlan <vlan>`
6	Set the priority on the remote management interface. `interface remote set {priority <NUMBER: 0..7>}` where priority <NUMBER: 0..7> is the 802.1D priority tag.
7	Configure the remote interface's IP address: `interface remote set {ip <IP address>}` where ip <IP address> sets the IPv4 address of the interface.

Copyright© 2022 Ciena® Corporation

8 Configure the global IP interface attributes.

```
interface remote set {gateway <IP address>}
```

where

gateway <IP address>	is the default IPv4 gateway.

9 Configure the l2-transform on the VC transform.

```
virtual-circuit transform create l2-transform <l2-
transform> {vid <NUMBER: 1..4094>} [pcp <NUMBER: 0..7>]
```

where

<l2-transform>	creates the transform profile.
{vid <NUMBER: 1..4094>}	is the VLAN ID.
[pcp <NUMBER: 0..7>]	is the priority code point.

10 Create the l2-transform on the VLAN.

```
virtual-circuit ethernet create vc <vc> {vlan <VLAN>} l2-
transform <l2-transform profile>
```

where

vc <vc>	is the virtual circuit name.
vlan <VLAN>	is the provider VLAN
l2-transform>	is the transform profile.

11 Create the virtual switch on the virtual circuit.

```
virtual-switch ethernet create vs <vs> vc <vc>
```

where

vs <vs>	is the virtual switch name.
vc <vc>	is the virtual circuit name.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example configures a remote management interface on a VC LAN on a VC transform. This configuration results in double-tagged management traffic with vlan 2000 being the outer tag with a p-bit 2 and the inner tag being 4094 with a p-bit of 4.

```
virtual-circuit transform create l2-transform PushVID2000 vid 2000 pcp2
dhcp client disable
dhcpv6 client disable
vlan create vlan 4094
vlan add vlan 4094 port 24
interface remote set vlan 4094
interface remote set priority 4
interface remote set ip 10.10.71.101/26
interface set gateway 10.10.71.65
virtual-circuit ethernet create vc vcVzMgmt vlan 4094 l2-transform Push
ViD2000
virtual-switch ethernet create vs vs VzMgmt vc vc VzMgmt
```


Copyright© 2022 Ciena® Corporation

Procedure 7-19
Configuring a virtual circuit transform to a virtual circuit

Configure a VC transform to a VC.

Step	Action
1	Assign a virtual circuit transform to a virtual circuit: `virtual-circuit ethernet set vc <vc-name> l2-transform <vc-transform-name>`

—end—

Example

This example assigns a virtual circuit transform to a virtual circuit:

```
virtual-circuit ethernet set vc vc1000 l2-transform ptag1000
```


Copyright© 2022 Ciena® Corporation

Procedure 7-20
Removing a virtual circuit transform from a virtual circuit

Remove a VC transform from a VC.

Step	Action
1	Remove a VC transform from a VC: `virtual-circuit ethernet unset vc <vc-name> l2-transform` **—end—**

Example

This example removes a VC transform from a VC.

```
virtual-circuit ethernet unset vc vc1000 l2-transform ptag1000
```


Copyright© 2022 Ciena® Corporation

Procedure 7-21
Setting the virtual circuit transform to default values

Set the VC transform to default values.

Step	Action
1	Set the VC transform to default values: `virtual-circuit transform unset l2-transform <VC transform name> {dei}{pcp}{tpid}` **—end—**

Example

This example sets the VC transform to default values.

```
virtual-circuit transform unset l2-transform ptag1000 dei pcp tpid
```


Copyright© 2022 Ciena® Corporation

Procedure 7-22
Enabling or disabling a virtual circuit transform

Enable or disable a VC transform.

Step	Action

To enable a VC transform

1 Enable a VC transform:

```
virtual-circuit transform enable l2-transform <VC
transform name>
```

To disable a VC transform

2 Disable a VC transform:

```
virtual-circuit transform disable l2-transform <VC
transform name>
```

—end—

Example

This example disables a VC transform.

```
virtual-circuit transform disable l2-transform ptag1000
```


Copyright© 2022 Ciena® Corporation

Procedure 7-23
Deleting a virtual circuit transform

Delete a VC transform.

Step	Action
1	Delete a VC transform: `virtual-circuit transform delete l2-transform <VC transform name>`

—end—

Example

This example deletes a VC transform.

```
virtual-circuit transform delete l2-transform ptag1000
```


Copyright© 2022 Ciena® Corporation

Procedure 7-24
Displaying virtual circuit transform

You can display a specific VC transform or all VC transforms.

Step	Action

To display a VC transform

1 Display a specific VC transform:

```
virtual-circuit transform show l2-transform <VC transform
name>
```

To display all VC transforms

2 Display all VC transforms:

```
virtual-circuit transform show
```

—end—

Example

This example shows the output from the virtual-circuit transform show command for the VC transform pvid1000:

```
> virtual-circuit transform show l2-transform pvid1000

+----------- L2-TRANSFORM pvid1000 INFO ------------+
| Parameter        | Value                           |
+------------------+---------------------------------+
| Status           | Enabled                         |
| TPID             | 88A8                            |
| PCP              | 5                               |
| DEI              | 0                               |
| VID              | 1000                            |
| VC Count         | 2                               |
+------------------+---------------------------------+
```

This example shows the output from the virtual-circuit transform show command:

```
> virtual-circuit transform show

+---------------------------- L2-TRANSFORM TABLE -----------------------------+
| Name                              | Status    | TPID | VID  | PCP | DEI | VCs  |
+-----------------------------------+-----------+------+------+-----+-----+------+
| pvid1000                          | Enabled   | 88A8 | 1000 |   5 |   0 |    2 |
+-----------------------------------+-----------+------+------+-----+-----+------+
```


Copyright© 2022 Ciena® Corporation

Procedure 7-25
Displaying system-reserved port information for the 5160 switch

Display system-reserved port information for the 5160 switch to confirm the configuration of the system-reserved port.

Step	Action

To display system-reserved port information on the 5160

1 Display the system-reserved port information on the 5160:

```
chassis system-rsvd-port show
```

To display system-reserved port information for the vc-transform feature on the 5160

2 Display the system-reserved port information for the vc-transform feature on the 5160 (optional):

```
chassis system-rsvd-port show feature vc-transform
```

—end—

Copyright© 2022 Ciena® Corporation

Example

This example displays output from the chassis system-rsvd-port show command with a reserved port:

```
5160> chassis system-rsvd-port show

+----------------- SYSTEM RESERVED PORT INFO -----------------+
| Port                | Feature                               |
+---------------------+---------------------------------------+
| 19                  | vc-transform                          |
+---------------------+---------------------------------------+
```

This example displays output from the chassis system-rsvd-port show feature vc-transform command:

```
5160> chassis system-rsvd-port show feature vc-transform

+----------------- SYSTEM RESERVED PORT INFO -----------------+
| Feature             | Port                                  |
+---------------------+---------------------------------------+
| vc-transform        | 19                                    |
+---------------------+---------------------------------------+
```

This example displays output from the chassis system-rsvd-port show command without a reserved port:

```
5160*> chassis system-rsvd-port show

+----------------- SYSTEM RESERVED PORT INFO -----------------+
| Port                | Feature                               |
+---------------------+---------------------------------------+
| No ports have been reserved for feature use                 |
+---------------------+---------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 7-26
Preventing NNI ingress frames without P-Tag from being switched

A virtual-switch or VS cross-connect configured with L2 VC transform services accepts and switches frames received on the NNI port with appropriate VIDs in the P-Tag and S-Tag. P-VID matching the VC transform configured VID and the S-VID matching the VLAN configured on the virtual circuit.

By default, the same virtual switch or VS cross-connect also accepts and switches frames received on the NNI port where the outer tag is the S-Tag whose VID that matches the VLAN configured on the virtual circuit.

This figure shows an example of a third-party network access node receiving P-Tag and S-Tag frames.

Figure 7-33
Third-party network access node

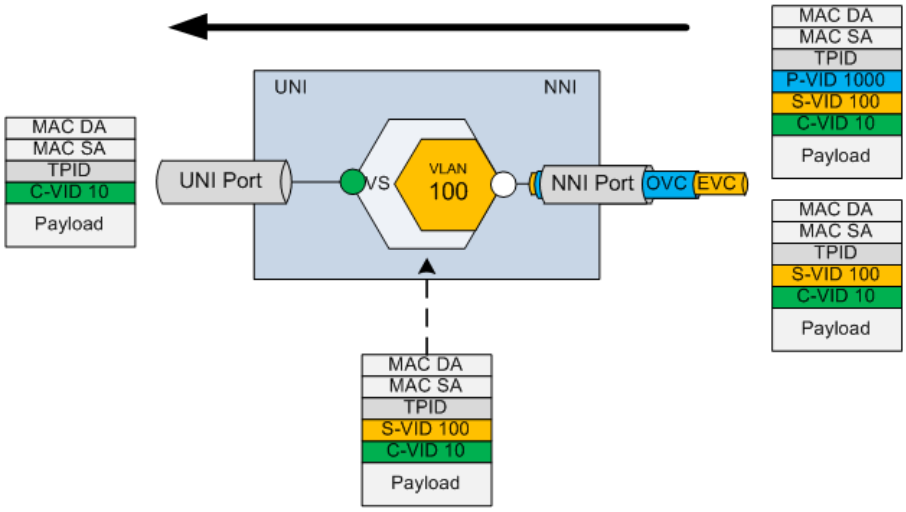

It is uncommon for an access node on a third-party network to receive frames with an outer tag other than the P-Tag. If the node receives an undesired frame on the NNI port with an outer S-Tag, those frames can be dropped by the node by configuring VLAN translation on the NNI port.

Copyright© 2022 Ciena® Corporation

Step	Action
1	Prevent the NNI ingress frames without P-Tag from being switched: `vlan translate add port <nni-port> vid <s-vid> vlan 4095` where port <nni-port> is the port(s) to add to the VLAN translation. vid <s-vid> is the VLAN ID. **—end—**

Example

This example prevents NNI ingress frames without P-Tag from being switched:

```
>vlan translate add port 1 vid 100 vlan 4095
```


Copyright© 2022 Ciena® Corporation

Procedure 7-27
Configuring L2 PFGs

Configure L2 PFGs to restrict forwarding of L2 traffic between groups of ports in a shared VLAN or virtual switch.

L2 PFG contains three levels of configuration:

- global administrative state, which by default is enabled
- PFG policy set specifications
- VLAN- and virtual switch-based configurations settings

The forwarding restrictions defined by L2 PFG become active, that is, operationally-enabled, when all three levels of configuration are enabled.

By default, L2 PFG 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© 2022 Ciena® Corporation

This table lists supported domains by platform.

Table 7-22
Supported domains by platform

Platform	Domain
3903	management
3903x	management
3904	management
3905	management
3906	management
3924	management, universal
3926	management, universal
3928	management, universal
3942	management
5142	management, universal
5160	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.

Copyright© 2022 Ciena® Corporation

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 <VlanIdList> [name <String[31]>]
vlan add vlan <VLAN ID> port <PORT1>
virtual-circuit ethernet create vc <vc> vlan <VLAN ID>
virtual-switch ethernet create vs <vs> vc <vc>
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—

Copyright© 2022 Ciena® Corporation

Example

The Sample configuration figure illustrates a simplified example of PFG use on a 39xx platform. Port 10 connects to a core network and ports 5 and 8 are connected to two 39xx switches. The 39xx devices represent two separate customer networks.

Figure 7-34
Sample configuration

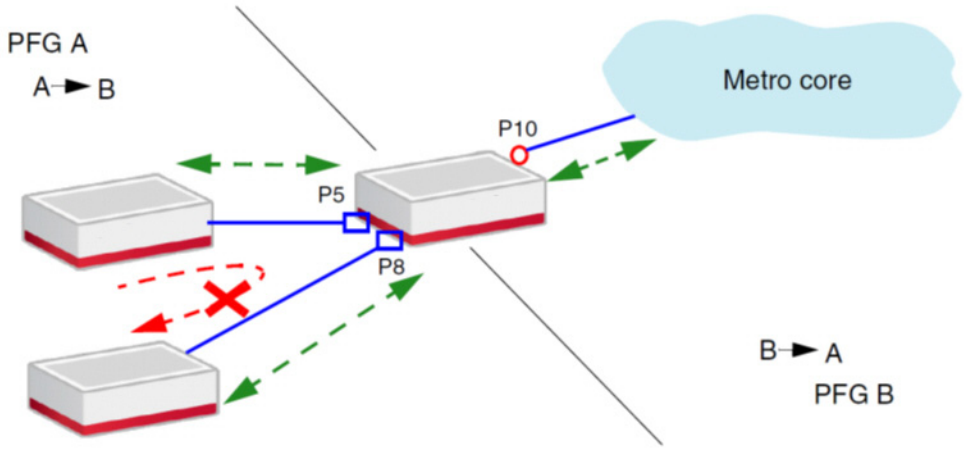

This example configures one service instance while provisioning two attachment customer networks. This configuration sets up a PFG such that the 39xxs can access the core network but cannot communicate with each other.

```
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
```


Copyright© 2022 Ciena® Corporation

Procedure 7-28
Configuring port-based PFGs

Configure port-based PFGs to alter the forwarding behavior of ports on all VLANs.

Port-based PFG contains four levels of configuration:

- global administrative state, which by default is enabled
- global port-based administrative state
- PFG egress profile configuration
- port PFG configuration

Global administrative state determines whether PFG is enabled globally. Disabling PFG removes the forwarding restrictions of PFG provisioned ports.

Global port-based administrative state determines whether port-based PFG is enabled globally. Disabling PFG globally removes the forwarding restrictions of PFG provisioned ports. Port-based PFG has its own administrative state. This is disabled if port-based PFG is disabled. PFG is enabled, port-based PFG may be enabled or disabled.

PFG egress profile configuration settings control the parameters that may be applied to each port. These settings are applied to all ports which share that egress profile.

Port PFG configuration settings assigns up to eight egress profiles for each port. That means one per flood traffic type combination.

Egress profiles must contain optional forwarding or blocking based on these data traffic types:

- unknown unicast
- unknown multicast
- broadcast
- none of these
- or any combination of these

By default, traffic type is all, meaning that all traffic is forwarded to the given egress ports or LAGs. For more information, refer to "Configuring egress profile with traffic types" on page 7-121.

Copyright© 2022 Ciena® Corporation

Step	Action

1	Create an egress profile for traffic egressing on ports 1, 2, 3 and 5 by one of these two methods:
	`private-forwarding-groups port-forwarding create egress-profile EPOL01`
	`private-forwarding-groups port-forwarding add egress-profile EPOL01 port 1-3,5`
	or
	`private-forwarding-groups port-forwarding create egress-profile EPOL01 port 1-3,5`

2	Assign this policy to port 8: `private-forwarding-groups port-forwarding add port 8 egress-profile EPOL01`

3	Enable port-based PFG:
	`private-forwarding-groups port-forwarding enable`
	`private-forwarding-groups enable`

—end—

Example

This figure shows a partial configuration in which data traffic ingressing on port 8 may be forwarded only to ports 1, 2, 3 and 5.

Figure 7-35
Port-based PFG sample partial configuration

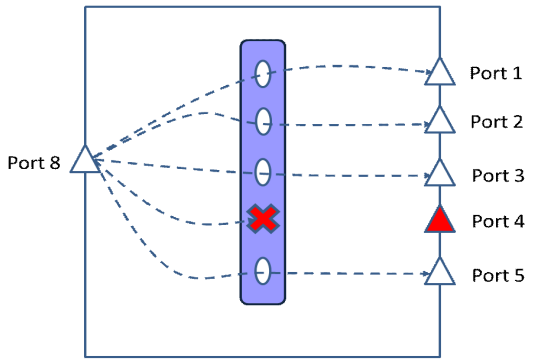

This example configures an egress profile for traffic egressing on ports 1,2, 3 and 5. The policy is set to port 8 and port-based PFG is enabled.

```
private-forwarding-groups port-forwarding create egress-profile EPOL01 port
1-3,5
private-forwarding-groups port-forwarding add port 8 egress-profile EPOL01
private-forwarding-groups enable
```


Copyright© 2022 Ciena® Corporation

```
private-forwarding-groups port-forwarding enable
```

This figure shows a configuration with 8 ports. All forwarding configuration is shown as unidirectional with bidirectional connections being achieved with unidirectional configuration in both directions.

Figure 7-36
Unidirectional port-based PFG sample configuration

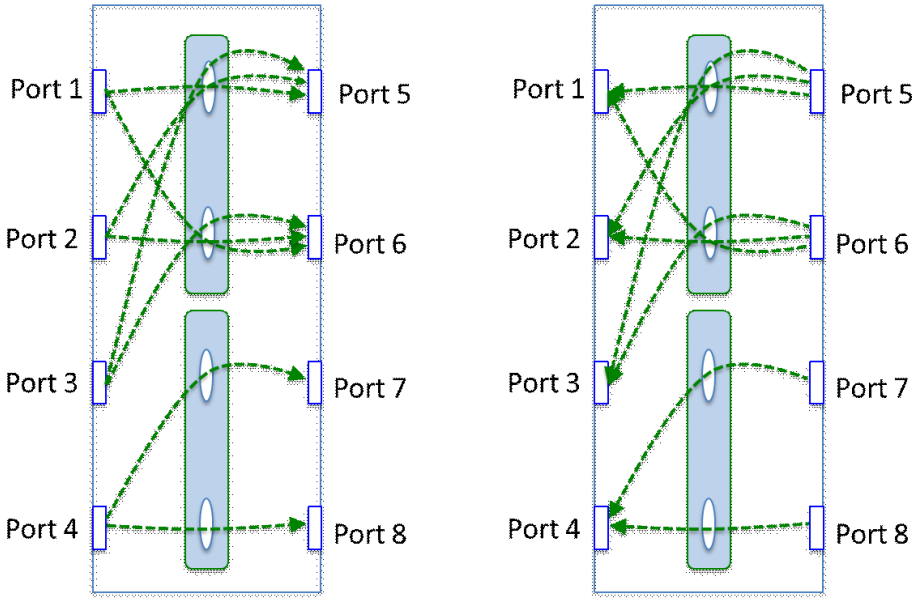

This example creates egress profiles for traffic egressing on ports 1, 2, 3, 4, 7 and 8. The policies are assigned to the desired port, and port-based PFG is enabled.

```
private-forwarding-groups port-forwarding create egress-profile EP_un01 port
5,6
private-forwarding-groups port-forwarding create egress profile EP_un02 port
7,8
private-forwarding groups port-forwarding create egress-profile EP_nu01 port
1-3
private-forwarding-groups port-forwarding create egress-profile EP_nu02 port
4
private-forwarding groups port-forwarding add port 1-3 egress-profile EP_un01
private-forwarding-groups port-forwarding add port 4 egress-profile EP_un02
private-forwarding-groups port-forwarding add port 5,6 egress-profile EP_nu-
01
private-forwarding-groups port-forwarding add port 7,8 egress-profile EP_un02
private-forwarding-groups port-forwarding enable
private-forwarding-groups enable
```


Copyright© 2022 Ciena® Corporation

This figure shows how bidirectional port forwarding may be configured. Any port which is in more than one egress profile applied to a port must not have conflicting flood traffic types. That is, a port may not be in more than one egress profile which is assigned to any port if its flood traffic types are not the same.

Figure 7-37
Bidirectional port-based PFG sample configuration

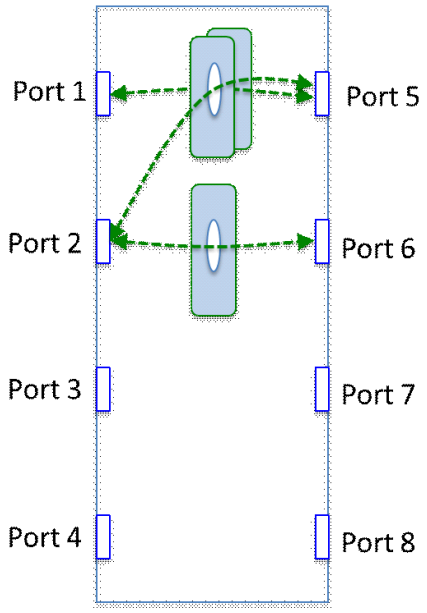

This example creates egress profiles for traffic egressing on ports 1, 2, 5 and 6. The policies are assigned to the desired port, and port-based PFG is enabled.

```
private-forwarding-groups port-forwarding create egress-profile EP_01 port
1,5
private-forwarding-groups port-forwarding create egress profile EP_02 port
2,5
private-forwarding groups port-forwarding create egress-profile EP_03 port
2,6
private-forwarding-groups port-forwarding add port 1 egress-profile EP_01
private-forwarding-groups port-forwarding add port 2 egress-profile
EP_02,EP_03
private-forwarding-groups port-forwarding add port 5 egress-profile
EP_01,EP_02
private-forwarding-groups port-forwarding add port 6 egress profile EP_03
private-forwarding-groups port-forwarding enable
private-forwarding-groups enable
```


Copyright© 2022 Ciena® Corporation

Procedure 7-29
Configuring egress profile with traffic types

The traffic types previously described may be added to each egress profile singly or in combinations of unknown-unicast, unknown multicast and broadcast. The egress profile can have unknown-unicast and unknown-multicast traffic with blocked broadcast traffic. The switch hardware cannot selectively block known-unicast or known multicast traffic. These can only be blocked by removing the respective egress ports from the egress profile.

There are eight combinations of allowed flood traffic types:

- none (no unknown-unicast, unknown-multicast or broadcast)
- unknown-unicast
- unknown-multicast
- broadcast
- unknown-unicast and unknown-multicast
- unknown-unicast and broadcast
- unknown-multicast and broadcast
- unknown-unicast, unknown-multicast and broadcast

The default is unknown-unicast, unknown-multicast and broadcast.

To allow a port to block broadcast traffic to one egress port, but allow it on another requires that the ports be permitted to have more than one egress profile assigned to them.

Step	Action

1	Create an egress profile for traffic egressing on ports 5 and 6:
	`private-forwarding-groups port-forwarding create egress-profile EP_un01 port 5`
	`private-forwarding-groups port-forwarding create egress-profile EP_un02 port 5 flood-type unknown-ucast, unknown-mcast`
	`private-forwarding-groups port-forwarding create egress-profile EP_un03 port 6 flood-type none`

2	Assign these policies to the desired ports:
	`private-forwarding-groups port-forwarding add port 1,2 egress-profile EP_un01,EP_un03`
	`private-forwarding-groups port-forwarding add port 3 egress-profile EP_un01,EP_un03`

Copyright© 2022 Ciena® Corporation

3 Enable port-based PFG:

```
private-forwarding-groups port-forwarding enable
private-forwarding-groups enable
```

—end—

Example

This figure shows a sample port-based PFG configuration with restrictions on flood traffic types for some ports.

Figure 7-38
Port-based PFT sample configuration with flood traffic types

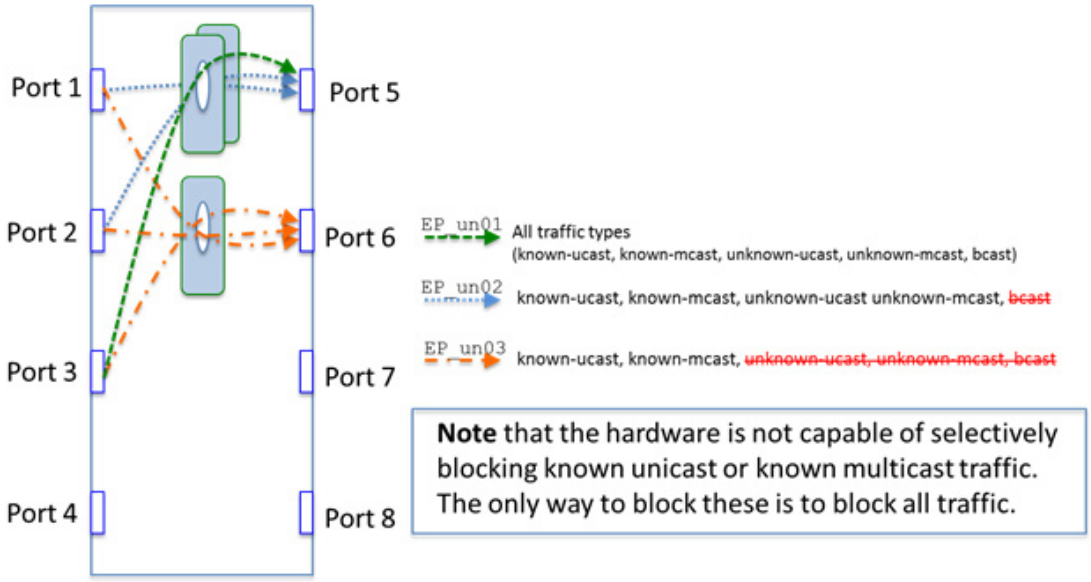

This example creates egress profiles for traffic egressing on ports 5 and 6. The policies are assigned to the desired port, and port-based PFG is enabled.

```
private-forwarding-groups port-forwarding create egress-profile EP_un01 port
5
private-forwarding-groups port-forwarding create egress-profile EP_un02 port
5 flood-type unknown-ucast,unknown-mcast
private-forwarding-groups port-forwarding create egress-profile EP_un03 port
6 flood-type none
private-forwarding-groups port-forwarding add port 1,2 egress-profile
EP_un01,EP_un03
private-forwarding-groups port-forwarding add port 3 egress-profile
EP_un01,EP_un03
private-forwarding-groups port-forwarding enable
private-forwarding-groups enable
```


Copyright© 2022 Ciena® Corporation

When two or more egress profiles are assigned to any port, each of the egress ports they contain must have a matching flood traffic type with the other egress profiles, as shown in this figure.

Figure 7-39
Port-based PFG sample configuration with flood traffic conflict

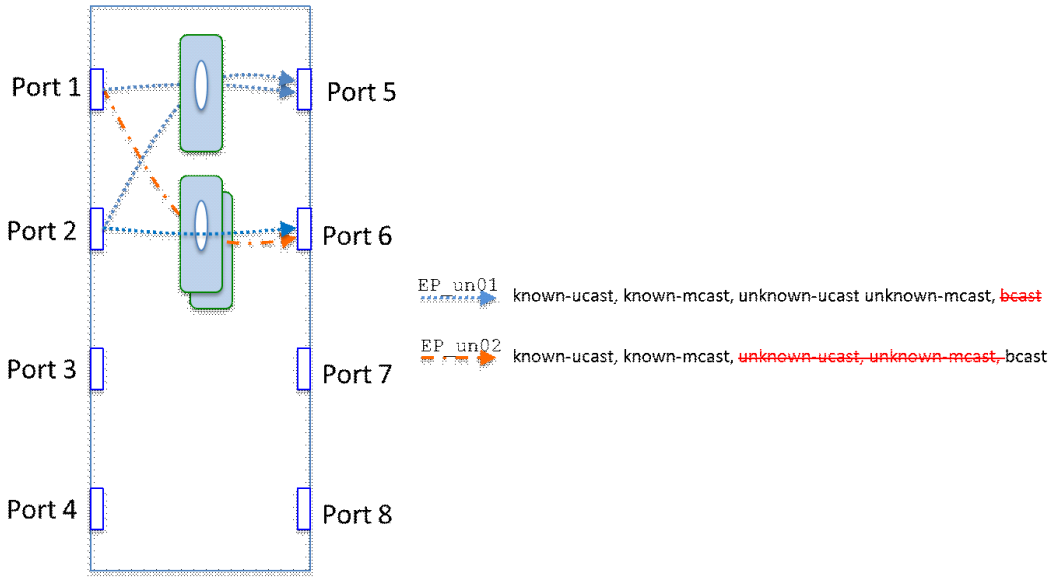

This example creates egress profiles for traffic egressing on ports 5 and 6. An attempt is made to assign policies to the desired ports. These have conflicting allow-flood types to egress port 6.

```
private-forwarding-groups port-forwarding create egress-profile EP_un01 port
5,6 flood-type unknown-ucast, unknown-mcast
private-forwarding-groups port forwarding create egress -profile EP_un02 port
6 flood-type bcast
private-forwarding-groups port-forwarding add port 1 egress-profile
EP_un01,EP_un02
private-forwarding-groups port-forwarding add port 2 egress-profile EP_un01
```


Copyright© 2022 Ciena® Corporation

Procedure 7-30
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`
5	Disable port-based PFG: `private-forwarding-groups port-forwarding disable`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 7-31
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

This example shows sample output for the virtual switch with an identifier of vs150.

```
> virtual-switch ethernet show vs vs150

+-------------------- QinQ/PBT VIRTUAL SWITCH INFO ---------------------+
| Parameter              | Value                                        |
+------------------------+----------------------------------------------+
| Name                   | vs150                                        |
| ID                     | 5                                            |
| Description            |                                              |
| Mode                   | default                                      |
| Active VLAN            | (150)   VLAN#150                             |
| Reserved VLAN          | (---)                                        |
| 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                                            |
| PFG State              | Disabled                                     |
|     Fwd-Policies       | A -> AB  B -> A                              |
|------------------------------------------------------------------------|
+-------------------------- Virtual Switch Members ----------------------+
| Port     | Subscriber | Translate | Encap Cos     |Encap Fixed | Fwd     | Stats |
|          | VLAN       | Tag       | Policy        | Dot1DPri   | Group   |       |
|---------+-----------+----------+-------------+-----------+--------+------+
| 10       | 10         | ---       | vs-inherit    | 2          | A       | off   |
+------------------------------------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 7-32
Displaying virtual switches

Display virtual switches to learn which reserved VLAN is associated with a virtual switch.

Step	Action
1	Display the reserved VLAN associated with virtual switches: `virtual-switch show`

—end—

Example

This example shows sample output for the virtual-switch show command.

```
> virtual-switch show

+------------ VIRTUAL SWITCH RESERVED VLANS -----------+
| VLAN                  | Virtual Switch [Attached]     |
+-----------------------+-------------------------------+
| 4000                  |                               |
| 4001                  |                               |
| 4002                  |                               |
| 4003                  | (1)   vs2000                  |
| 4004                  |                               |
| 4005                  |                               |
| 4006                  |                               |
| 4007                  |                               |
| 4008                  |                               |
| 4009                  |                               |
| 4010                  |                               |
+-----------------------+-------------------------------+

+------------ MPLS VIRTUAL SWITCH TABLE ------------------------------+
|                  |             | VCircuit  | ACircuit |              |
| Name             | CPT | Mode  | Count     | Count    | Description  |
+------------------+-------------+-----------+----------+--------------+
| No entries found                                                    |
+------------------+-------------+-----------+----------+--------------+

+-------------------- QinQ/PBT VIRTUAL SWITCH TABLE ------------------------+
|               |  VLAN  |      |                 | Subscriber|              |
| Name          | Active | CPT  | Virtual Circuit |   Count   | Description  |
+---------------+--------+------+-----------------+-----------+--------------+
| newV3         | 1200   | Dis  | New VC          | 0         |              |
+---------------+--------+------+-----------------+-----------+--------------+
```


Copyright© 2022 Ciena® Corporation

This example shows sample output for the virtual switch VS1.

```
> virtual-switch show vs VS1

+-------------------- QinQ/PBT VIRTUAL SWITCH INFO ----------------------+
| Parameter             | Value                                          |
+-----------------------+------------------------------------------------+
| Name                  | VS1                                            |

| ID                    | 1                                              |
| Description           |                                                |
| Mode                  | default                                        |
| Active VLAN           | (500)   VLAN#500                               |
| Reserved VLAN         | (---)                                          |
| CPT Status            | Disabled                                       |
| CPT Method            | l2pt                                           |
| Transparent Validate  | off                                            |
| Mac Learning          | Enabled                                        |
| VC                    | (1)     VC1                                    |
| Subscr Dot1D Policy   | leave                                          |
| Decap CoS Policy      | leave                                          |
| Decap Fixed Dot1DPri  | 2                                              |
| Encap CoS Policy      | fixed                                          |
| Encap Fixed Dot1DPri  | 2                                              |
| PFG State             | Disabled                                       |
|     Fwd-Policies      | A -> AB  B -> A                                |
|-----------------------------------------------------------------------|

+--------------------------- Virtual Switch Members ----------------------------+
| Port     | Subscriber | Translate | Encap Cos     |Encap Fixed | Fwd     | Stats |
|          | VLAN       | Tag       | Policy        | Dot1DPri   | Group   |       |
|---------+------------+-----------+--------------+------------+--------+-------+
| 10       | 10         | ---       | vs-inherit    | 2          | A       | off   |
+-------------------------------------------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 7-33
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
- port-based PFG settings
- port-based PFG egress profiles
- port-based PFG settings for a port

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`
To display port-based PFG settings	
6	Display port-based PFG settings: `private-forwarding-groups port-forwarding show`

Copyright© 2022 Ciena® Corporation

To display port-based PFG egress profiles

7 Display port-based PFG egress profiles:

```
private-forwarding-groups port-forwarding egress-
[profile <egressProfile>]
```

where

show [profile <egressProfile>]	is the PFG egress profile.

To display port-based PFG settings for a port

8 Display port-based settings for a port:

```
private-forwarding-groups port-forwarding show port
<port>
```

where

show <port>	is the port.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 7-34
Configuring EPL and EVPL for E-Access service types

Configure User Premises Equipment (UPE) and Network Premises Equipment (NPE) to support E-Access service types.

Step	Action

To configure the UPE as an EPL or EVPL

1 Configure the UNI and NNI ports maximum frame size.

```
port set port <PortList> max-frame-size <NUMBER: 1522-9216>
```

2 Create a provider VLAN (SVID).

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

3 Configure the provider VLAN TPID.

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

4 Add the NNI port to the provider VLAN.

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

5 Create a management VLAN.

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

6 Add the management port to the remote management VLAN.

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

7 Remove the ports from the default management VLANs.

```
vlan remove vlan <VlanList> port <PortNameList>
```

8 Change the remote interface to the new remote management VLAN.

```
interface remote set vlan <Vlan>
```

9 Configure the PVID and VLAN for egress untagged traffic.

```
port set port <PortList> egress-untag-vlan <Vlan>
```


Copyright© 2022 Ciena® Corporation

10 Configure the acceptable frame type for the UNI port.

```
port set port <port> acceptable-frame-type <all|tagged-only|untagged-only>
```

where

port <port>	is the port(s) to set.
[acceptable-frame-type <all|tagged-only|untagged-only>]	sets the acceptable frame types mode as follows: • all - all frames are allowed to ingress, regardless of their tag status. • tagged-only - only tagged frames are allowed to ingress. • untagged-only - only untagged frames are allowed to ingress. ***Note:*** When acceptable frame types are configured on a non-virtual switch UNI, untagged0only mode filters out all three common TPIDs of 0x8100, 0x8898, and 0x9100. For a virtual switch UNI, only the TPID of 0x8100 is recognized as a tagged frame.

11 Create an Ethernet virtual circuit using the provider VLAN.

```
virtual-circuit ethernet create {vc <vc>} {vlan <Vlan>} [l2-transform <vc-transform-name>]
```

12 Create Ethernet virtual switches using the Ethernet virtual circuit.

```
virtual-switch ethernet create vs <vs> vc <vc>
```

13 Disable the LLDP on the UNI port.

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

14 Add UNI PORT+CVID members to the Ethernet 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>] [pcp <NUMBER: 0-7>]
```

Note: When the vlan <vlan list> parameter is specified, the UPE is configured for an EVPL. When the vlan <vlan list> parameter is left out, the UPE is configured for an EPL.

15 Set the UNI and NNI port VLAN Ethertype value and policy.

```
virtual-circuit ethernet set port <PortList> vlan-ethertype <8100|9100|88A8> vlan-ethertype-policy <all|encap-only|vlan-tpid>
```

To configure the NPE

16 Configure the UNI and NNI ports maximum frame size.

```
port set port <PortList> max-frame-size <NUMBER: 1522-9216>
```


Copyright© 2022 Ciena® Corporation

17 Create a provider VLAN (SVID).

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

18 Configure the provider VLAN TPID.

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

19 Add the UNI and NNI ports to the provider VLAN.

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

20 Create a management VLAN.

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

21 Add the management port to the remote management VLAN.

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

22 Remove the ports from the default management VLANs.

```
vlan remove vlan <VlanList> port <PortNameList>
```

23 Change the remote interface to the new remote management VLAN.

```
interface remote set vlan <Vlan>
```

24 Configure the PVID and VLAN on the management port for egress untagged traffic.

```
port set port <PortList> egress-untag-vlan <Vlan>
```

25 Configure the acceptable frame type for the ENNI port.

```
port set port <PortList> acceptable-frame-type
<all|tagged-only|untagged-only>
```

—end—

Example

This figure shows an example of an E-Access EPL/EVPL where the ENNI receives double tagged traffic with the outer SVID tag of 200 and inner CVID tag of 100.

Figure 7-40
E-Access EPL/EVPL

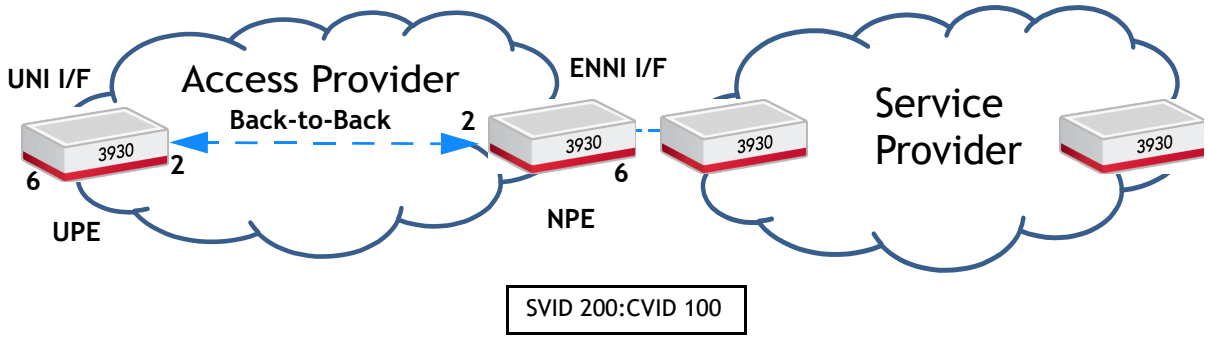

Copyright© 2022 Ciena® Corporation

Configure the UPE

```
port set port 2 max-frame-size 9216
port set port 6 max-frame-size 9216
vlan create vlan 200 name S-tag
vlan set vlan 200 egress-tpid 88A8
vlan add vlan 200 port 2
vlan create vlan 128
vlan add vlan 128 port 1
vlan remove vlan 1 port 1
vlan remove vlan 1 port 2
vlan remove vlan 127 port 6
interface remote set vlan 128
port set port 1 pvid 128 egress-untag-vlan 128
port set port 6 acceptable-frame-type all
virtual-circuit ethernet create vc E-Access vlan 200
virtual-switch ethernet create vs E-Access vc E-Access
lldp set port 6 disable
virtual-switch ethernet add vs E-Access port 6 vlan 100
virtual-circuit ethernet set port 2 vlan-ethertype 88A8 vlan-ethertype-policy
vlan-tpid
virtual-circuit ethernet set port 6 vlan-ethertype 88A8 vlan-ethertype-policy
vlan-tpid
```

Configure the NPE:

```
port set port 2 max-frame-size 9216
port set port 6 max-frame-size 9216
vlan create vlan 200 name S-tag
vlan set vlan 200 egress-tpid 88A8
vlan add vlan 200 port 2
vlan add vlan 200 port 6
vlan create vlan 128
vlan add vlan 128 port 1
interface remote set vlan 128
port set port 1 pvid 128 egress-untag-vlan 128
port set port 6 acceptable-frame-type all
```


Copyright© 2022 Ciena® Corporation

Procedure 7-35
Configuring CFM for E-Access service types

Configure CFM over UPE and NPE for E-Access service types created in Procedure 7-34.

Step	Action
To configure the UPE	
1	Enable CFM. `cfm enable`
2	Create a CFM service for the E-access virtual switch. `cfm service create`
3	Enable the CFM service. `cfm service enable`
4	Create an up MEP for the service for the UNI PORT+CVID. `cfm mep create`
5	Send DMMs to the remote MEP: `cfm delay send`
6	Set CFM service LMM interval: `cfm service set`
7	Send LMMs to a remote MEP: `cfm frame-loss send`
8	Display the MEP. `cfm mep show`
9	Display the remote MEP `cfm remote-mep show`
To configure the NPE	
10	Enable CFM. `cfm enable`
11	Turn off 802.1ad strict mode. `cfm set dot1ad-strict off`
12	Create a CFM service for the E-access virtual switch. `cfm service create`
13	Enable the CFM service. `cfm service enable`
14	Create an up MEP for the service for the UNI PORT+CVID. `cfm mep create`

Copyright© 2022 Ciena® Corporation

15 Set the MEP attributes to use the tag-vid for the CVID and turn on the tag-mode.

```
cfm mep set
```

16 Send DMMs to the remote MEP:

```
cfm delay send
```

17 Set the CFM service LMM interval:

```
cfm service set
```

18 Send LMMs to a remote MEP:

```
cfm frame-loss send
```

19 Display the MEP.

```
cfm mep show
```

20 Display the remote MEP

```
cfm remote-mep show
```

—end—

Example

This example configures

- "UPE CFM for EVPL" on page 7-135
- "NPE CFM for EVPL" on page 7-136
- "UPE CFM for EPL" on page 7-136
- "NPE CFM for EPL" on page 7-136

UPE CFM for EVPL

```
cfm enable
cfm service create vs E-Access name ENNI-test md md4 next-mepid 611
cfm service enable service ENNI-test
cfm mep create service ENNI-test port 6 vlan 100 type up mepid 611
cfm delay send service ENNI-test local-mepid 611 mepid 711 iterate 0 repeat-
delay 1 delay-threshold 700 jitter-threshold 10
cfm service set service ENNI-test lmm-interval 1sec
cfm frame-loss send service ENNI-test local-mepid 611 mepid 711 iterate 0
repeat-delay 1
cfm mep show
cfm remote-mep show
```


Copyright© 2022 Ciena® Corporation

NPE CFM for EVPL

```
cfm enable
cfm set dot1ad-strict off
cfm service create vlan 200 name ENNI-test md md4 next-mepid 711
cfm service enable service ENNI-test
cfm mep create service ENNI-test port 6 type up mepid 711
cfm mep set service ENNI-test local-mepid 711 tag-vid 100
cfm mep set service ENNI-test local-mepid 711 tag-mode on
cfm delay send service ENNI-test local-mepid 711 mepid 611 iterate 0 repeat-
delay 1 delay-threshold 700 jitter-threshold 10
cfm service set service ENNI-test lmm-interval 1sec
cfm frame-loss send service ENNI-test local-mepid 711 mepid 611 iterate 0
repeat-delay 1
cfm mep show
cfm remote-mep show
```

UPE CFM for EPL

```
cfm enable
cfm service create vs E-Access name ENNI-test md md4 next-mepid 611
cfm service enable service ENNI-test
cfm mep create service ENNI-test port 6 type up mepid 611
cfm delay send service ENNI-test local-mepid 611 mepid 711 iterate 0 repeat-
delay 1 delay-threshold 700 jitter-threshold 10
cfm service set service ENNI-test lmm-interval 1sec
cfm frame-loss send service ENNI-test local-mepid 611 mepid 711 iterate 0
repeat-delay 1
cfm mep show
cfm remote-mep show
```

NPE CFM for EPL

```
cfm enable
cfm set dot1ad-strict on
cfm service create vlan 200 name ENNI-test md md4 next-mepid 711
cfm service enable service ENNI-test
cfm mep create service ENNI-test port 6 type up mepid 711
cfm delay send service ENNI-test local-mepid 711 mepid 611 iterate 0 repeat-
delay 1 delay-threshold 700 jitter-threshold 10
cfm service set service ENNI-test lmm-interval 1sec
cfm frame-loss send service ENNI-test local-mepid 711 mepid 611 iterate 0
repeat-delay 1
cfm mep show
cfm remote-mep show
```


Copyright© 2022 Ciena® Corporation

Procedure 7-36
Configuring HIM for E-Access service types

Configure hierarchical ingress metering (HIM) with traffic profiling on the UPE and NPE for E-Access service types created in Procedure 7-34. The configuration on the UNI and ENNI is the same for EVPL and EPL.

For details about HIM, refer to "Quality of Service configuration" on page 8-1.

Step	Action

To configure the UPE UNI

1 Set the traffic profiling meter provisioning to EIR.

```
traffic-profiling set meter-provisioning eir
```

2 Set the UNI port traffic profiling to advanced mode.

```
traffic-profiling set port <PortName> mode advanced
```

3 Create a standard profile associated with the UNI port.

```
traffic-profiling standard-profile create port <PortName>
```

4 Enable traffic profiling on the UNI port.

```
traffic-profiling enable port <PortName>
```

5 Enable traffic profiling globally:

```
traffic-profiling enable
```

6 Display the traffic profiling throughput statistics:

```
traffic-profiling standard- throughput port <PortName>
```

To configure the NPE ENNI

7 Set the ENNI port traffic profiling to advanced hierarchical-vlan.

```
traffic-profiling set port <PortName> mode hierarchical-
vlan
```

8 Create a standard profile parent and child profiles associated with the ENNI port and SVID.

```
traffic-profiling standard-profile create port <PortName>
```

9 Enable traffic profiling on the ENNI port.

```
traffic-profiling enable port <PortName>
```

10 Enable traffic profiling globally:

```
traffic-profiling enable
```

11 Display the traffic profiling statistics:

```
traffic-profiling standard- throughput port <PortName>
```


Copyright© 2022 Ciena® Corporation

12 Display the traffic profiling throughput statistics for the child profile(s):

```
traffic-profiling standard- throughput port <PortName>
```

—end—

Example

This example shows configuration to limit the traffic rate to 50MB on the UNI and apply a 3 Rate HIM on the ENNI.

UPE metering on UNI

```
traffic-profiling set meter-provisioning eir
traffic-profiling set port 6 mode advanced
traffic-profiling standard-profile create port 6 profile 1 name E-Access cir
50048 pir 50048 cbs 256 ebs 0 vs E-Access
traffic-profiling enable port 6
traffic-profiling enable
traffic-profiling standard- throughput port 6 traffic-profile E-access
```

NPE metering on ENNI

```
traffic-profiling set port 6 mode hierarchical-vlan
traffic-profiling standard-profile create port 6 profile 1 name flow1 cir
300224 pir 450048 cbs 16 ebs 16 statistics off vlan 200 child-mode standard-
dot1dpri
traffic-profiling standard-profile create port 6 profile 2 name cos0 cir
100032 pir 150016 cbs 128 ebs 128 dot1dpri 0 parent flow1
traffic-profiling standard-profile create port 6 profile 3 name cos3 cir
100032 pir 150016 cbs 128 ebs 128 dot1dpri 3 parent flow1
traffic-profiling standard-profile create port 6 profile 4 name cos5 cir
100032 pir 150016 cbs 128 ebs 128 dot1dpri 5 parent flow1
traffic-profiling enable port 6
traffic-profiling enable

traffic-profiling standard- statistics

traffic-profiling standard- throughput port 6 traffic-profile cos0
```


Copyright© 2022 Ciena® Corporation

Procedure 7-37
Configuring L2 Control Frame Tunneling for E-Access service types

Configure L2 control frame tunneling for E-Access service types created in Procedure 7-34. The configuration is the same for EVPL and EPL and on the UPE only.

For details about L2 control frame tunneling, refer to "Layer 2 Control Frame Tunneling configuration" on page 11-1.

Step	Action
1	Configure the port for untagged data and control traffic. `port set port <PortList> untagged-data-vs <VirtualSwitchName> untagged-ctrl-vs <VirtualSwitchName>`
2	Configure the acceptable frame type for the UNI port to all. `port set port <PortList> acceptable-frame-type all`
3	Enable L2 control frame tunneling on the port. `l2-cft enable port <PortList>`
4	Configure the protocols with l2-cft profiles. `l2-cft protocol add profile <L2-CFT Profile Name> ctrl-protocol <layer 2 control protocol> untagged-disposition <Untagged l2cp protocol processing disposition> tagged-disposition <Tagged l2cp protocol processing disposition>`

—end—

Example

This example shows configuration to forward LLDP.

```
vlan add vlan 300 port y
virtual-circuit ethernet create vc vc1 vlan 300
virtual-switch ethernet create vs vs1 vc vc1
virtual-switch ethernet add vs vs1 port x vlan 100
l2-cft create profile LLDP_Tunnel
l2-cft protocol add profile LLDP_Tunnel ctrl-protocol lldp untagged-
disposition forward
l2-cft set port x profile LLDP_Tunnel
port set port x untagged-ctrl-vs vs1
l2-cft enble port x
```


Copyright© 2022 Ciena® Corporation

Procedure 7-38
Configuring RFC 2544 Benchmarking for E-Access

Configure RFC 2544 Benchmarking for E-Access service types as created in Procedure 7-34.

For details about RFC 2544 Benchmarking, refer to “System performance testing with RFC 2544 benchmarking” in *39XX/51XX Switches and Platforms Advanced OAM Configuration.*

Step	Action

To configure the generator UPE

1 Set the port to test and mode.

```
benchmark create generator name <genName> port <port>
emix-h-frame-size <INT 64..9216> and mode <in-service |
out-of-service | vid-out-of-service>
```

where

generator name <genName>	is the generator name.
port <port>	is the port.
emix-h-frame-size <INT 64..9216>	indicates the emix frame size.
mode <in-service | out-of-service | vid-out-of-service>	indicates the mode of the test port, which is one of in-service, out-of-service, or vid-out-of-service. The default is in-service.

2 Create the benchmark generator profile.

```
benchmark profile create name <String[45]>
```

where

name <String[45]>	is the user-defined name for new profile.

Copyright© 2022 Ciena® Corporation

3 Set benchmark generator profile interval and duration attributes:

```
benchmark configuration set name <name> [interval <15Min
| 1Hr | 6Hr | Completion>] [duration <15Min | 1Hr | 6Hr |
24Hr | Indefinite | Once>]
```

where

name <name>	is the profile to configure.
interval <15Min | 1Hr | 6Hr | Completion>	is the interval time.
duration <15Min | 1Hr | 6Hr | 24Hr | Indefinite | Once>	is the duration time.

4 Set benchmark generator profile payload attributes:

```
benchmark profile payload set name <name> [dstmac <MAC
address: XX:XX:XX:XX:XX:XX>] [vlan-encap-type <untagged |
dot1q>] [c-vid <NUMBER: 1-4094>] [c-pcp <NUMBER>] [c-
color <green|yellow|green-yellow|green-yellow-red>]
[tpid <0x8100|0x9100|0x8x88A8>] [pdu-type <ethernet | ip
| udp-echo>] [custom-payload <String>] [ip-src-addr <IP
address>] {ip-dest-addr <IP address>] [ip-dscp <NUMBER:
0-63>]
```

where

name <name>	is the profile to configure.
dstmac <MAC address: XX:XX:XX:XX:XX:XX>	is the destination MAC address.
vlan-encap-type <untagged | dot1q>	is the VLAN encapsulation type. The default is untagged.
[c-vid<NUMBER: 1-4094>]	is the customer VLAN ID value.
[c-pcp <NUMBER>]	is the customer priority code point (PCP value)
[c-color <green | yellow | green-yellow | green-yellow-red>]	is the color of test traffic. The default is green. When green is selected, one stream is sent out with the DEI bit set to 0. With yellow only, one stream of test traffic is sent out with DEI set to 1. For green-yellow or green-yellow-red, two streams of test traffic are sent out. The green traffic is sent using bandwidth. Yellow is sent at excess bandwidth. Red is sent at excess-bandwidth * 1.25.
[tpid <0x8100 | 0x9100 | 0x88A8>]	is the customer tag protocol identifier (TPID) value. The default is 0x8100.

Copyright© 2022 Ciena® Corporation

where

pdu-type <ethernet | ip | udp-echo>]	is the test PDU type. The default is ethernet.
[custom-payload <String>]	is the custom payload to use with the selected PDU type.
ip-src-addr <IP address>]	is the IP source address.
[ip-dest-addr <IP address>]	is the IP destination address.
[ip-dscp <NUMBER: 0-63>	is the IP DSCP value.

5 Enable the benchmark profile:

```
benchmark test create name <testInstName> profile
<profInstName> [vtag-stack <cvlan>] [dst-mac <DSTMAC>]
```

```
benchmark test enable name <testInstName>
```

where

name <testInstName>	is the profile to enable.
profile <profInstName>	is the profile instance name.
vtag-stack <cvlan>	if vtag is specified, it overwrites the selected profile configured c-vlan. If dst-mac is specified, it overwrites the selected profile configured dst-mac.
dst-mac <DSTMAC>	is the destination MAC.

6 Enable the generator:

```
benchmark generator name <genName> enable
```

7 Enable the global state:

```
benchmark enable
```

8 Confirm the global and generator state and configuration:

```
benchmark show
```

9 Check connectivity between the generator and UNI reflector:

```
benchmark test trigger name <testInstName> connectivity-
test
```

where

name <testInstName>	is the test instance to test the connectivity for.
connectivity-test	triggers a connectivity test on the selected profile.

Copyright© 2022 Ciena® Corporation

10 Start the test:

```
benchmark test start name <testInstName>
```

To configure the reflector UPE

11 Set the mode.

```
benchmark reflector set name <refName> [mode <in-service
| out-of-service | vid-out-of-service>] [vendor-type
<ciena|other>][mac-validation <on|off>]
```

where

name <refName>	is the name of the reflector instance to set.
mode <in-service| out-of-service| vid-out-of-service>	indicates the mode of the test port, which is one of in-service, out-of-service, or vid-out-of-service. The default is in-service.
vendor-type <ciena | other>	sets the vendor-type as Ciena or other.
mac-validation <on | off>	enables or disables mac-validation.

Note: If you wish to switch from generator to reflector or vis-a-versa, you have to delete the benchmark instance and re-create it. To delete a benchmark instance/entity, it has to be disabled and all associated test instances must be deleted.

12 Set the reflector VID.

```
benchmark test create name <testInstName> {{vtag-stack
<vlan>} | vlan-untagged-data}
```

where

name <testInstName>	is the name of the test instance to create.
vtag-stack <vlan>	required by the reflector test instance. You can also use * as the vtag-stack to specify any equivalent to SAOS 6.15.1 benchmark reflect unset vid. You must specify either a vtag-stack or vlan-untagged-data. one of the two must be specified when in the reflector role.
vlan-untagged-data	generates or reflects untagged packets.

13 Enable the reflector.

```
benchmark reflector enable name <refName>
```

14 Enable the global state.

```
benchmark enable
```


Copyright© 2022 Ciena® Corporation

15 Confirm the configuration settings and that the global and reflector states are enabled.

```
benchmark show
```

—end—

Example

This example shows configuration to forward LACP.

NPE RFC 2544 reflector for EVPL

```
benchmark enable
benchmark create reflector name CustRef mode out-of-service port 2
benchmark reflector enable name CustRef
benchmark test create name 2544test vtag-stack 400
benchmark test enable name 2544 test
```

UPE RFC 2544 generator for EVPL

```
benchmark enable
benchmark create generator name CustGen port 1 mode in-service
benchmark generator enable name CustGen
benchmark create name 2544test profile 2544test
benchmark profile create name 2544test
benchmark profile configuration set name 2544test interval Completion
benchmark profile configuration set name 2544test duration 15Min
benchmark profile configuration set name 2544test bandwidth 500
benchmark profile configuration set name 2544test vid-validation off
benchmark profile payload set name 2544test dst-mac 00:23:8a:34:10:10
benchmark profile payload set name 2544test vlan-encap-type dot1q
benchmark profile payload set name 2544test c-vid 400
benchmark profile payload set name 2544test c-pcp 100
benchmark profile payload set name 2544test c-color green
benchmark profile payload set name 2544test tpid 0x88A8
benchmark profile payload set name 2544test pdu-type udp-echo
benchmark profile payload set name 2544test custom-payload
"\xFFABC\x01def\x00heLLO!"
benchmark profile payload set name 2544test
benchmark profile payload set name 2544test ip-src-add 198.18.1.16 ip-dest-
addr 198.18.2.16
benchmark profile payload set name 2544test ip-dscp 50
benchmark test create name 2544test profile 2544test
benchmark test enable name 2544test
benchmark test trigger name 2544test connectivity-test
benchmark test start 2544test
```


Copyright© 2022 Ciena® Corporation

Procedure 7-39
Configuring ENNI hairpin switching using sub-ports

Sub-port interfaces support ENNI hairpin switching. This procedure describes how to configure a virtual switch for hairpin switching between sub-port interfaces. When you create sub-port interfaces, you can specify whether you want to turn on statistics. This procedure is an example of MEF 51 E-transit Service.

For more information on links between sub-ports and MEF OVCs, refer to "Q-in-Q encapsulation" on page 7-2.

Step	Action

To configure ENNI hairpin switching using a tagged sub-port on the INNI

1 Create a virtual switch:

```
virtual-switch add reserved-vlan <VLAN ID>
virtual-switch ethernet create vs <vs>
```

2 Create a tagged INNI sub-port.

```
sub-port create sub-port <sub-port-name> parent-port
<port-name> vtag-stack <VLAN-ID> statistics <on|off>
```

Note: Optionally, statistics can be enabled on a per-sub-port basis. If you do not specify an option, the default is off.

3 Attach the INNI tagged sub-port interface to the virtual switch:

```
virtual-switch ethernet attach sub-port <sub-port-name>
vs <virtual-switch-name>
```

4 Create the sub-port interfaces for the tagged ENNI port:

```
sub-port create sub-port <sub-port-name> parent-port
<port-name> vtag-stack <VLAN ID> statistics <on|off>
```

5 Attach the tagged ENNI sub-port interfaces to the virtual switch

```
virtual-switch ethernet attach sub-port <sub-port-name>
vs <virtual-switch-name>
```

6 View your configuration.

```
sub-port show
```

To configure ENNI hairpin switching using an untagged sub-port on the INNI

7 Create a virtual switch:

```
virtual-switch add reserved-vlan <VLAN ID>
virtual-switch ethernet create vs <vs>
```


Copyright© 2022 Ciena® Corporation

8 Create an untagged INNI port.

```
sub-port create sub-port <sub-port-name> parent-port
<port-name> vlan-untagged-data statistics <on|off>
```

Note: Optionally, statistics can be enabled on a per-sub-port basis. If you do not specify an option, the default is off.

9 Attach the INNI untagged sub-port interface to the virtual switch:

```
virtual-switch ethernet attach sub-port <sub-port-name>
vs <virtual-switch-name>
```

10 Create the sub-port interfaces for the tagged ENNI port:

```
sub-port create sub-port <sub-port-name> parent-port
<port-name> vtag-stack <VLAN ID> statistics <on|off>
```

11 Attach the tagged ENNI sub-port interfaces to the virtual switch

```
virtual-switch ethernet attach sub-port <sub-port-name>
vs <virtual-switch-name>
```

12 View your configuration.

```
sub-port show
```

where

<VLAN ID>	is a particular VLAN ID which has not already been configured on the device.
<virtual-switch-name>	is a unique virtual-switch name to identify this virtual switch instance.
<sub-port-name>	is a unique sub-port name to identify this sub-port instance.
<port-name>	identifies a particular port configured on the device.
<on|off>	identifies whether statistics are on or off for a particular sub-port. The default is off.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example configures ENNI hairpin switching using a tagged INNI sub-port.

```
virtual-switch add reserved-vlan 4000
virtual-switch ethernet create vs vs

sub-port create sub-port inni.100 parent-port 1 vtag-stack 100 statistics on

virtual-switch ethernet attach sub-port inni.100 vs vs4000

sub-port create sub-port enni.2023 parent-port 2 vtag-stack 2023 statistics on
sub-port create sub-port enni.1028 parent-port 2 vtag-stack 1028 statistics on

virtual-switch ethernet attach sub-port enni.2023 vs vs4000
virtual-switch ethernet attach sub-port enni.1028 vs vs4000
```

This example configures ENNI hairpin switching using an untagged INNI sub-port.

```
virtual-switch add reserved-vlan 4000
virtual-switch ethernet create vs vs

sub-port create sub-port inni.untagged parent-port 1 vlan-untagged-data
statistics on

virtual-switch ethernet attach sub-port inni.untagged vs vs4000

sub-port create sub-port enni.1028 parent-port 2 vtag-stack 1028 statistics on
sub-port create sub-port enni.2023 parent-port 2 vtag-stack 2023 statistics on


virtual-switch ethernet attach sub-port enni.2023 vs vs4000
virtual-switch ethernet attach sub-port enni.1028 vs vs4000
```


Copyright© 2022 Ciena® Corporation

Procedure 7-40
Displaying statistics for sub-port interfaces

You can display the administrative and operational attributes for all sub-port interfaces on a device or for a specific list of sub-port interfaces. When the statistics option is specified, the statistics counts for all sub-port interfaces that have statistics currently enabled are displayed.

Step	Action

To display statistics for sub-ports with statistics enabled and non-zero

1 Display statistics for sub-port interfaces on the device.

```
sub-port show statistics
```

To display statistics for sub-ports with statistics enabled and non-zero with a specified parent port

2 Display statistics for sub-port interfaces that have the port specified as its parent port.

```
sub-port show parent-port <port-name> statistics
```

where

<port-name>	is the specified parent port you want to display statistics for.

To display statistics for sub-ports or a list of sub-ports

3 Display statistics for a specific sub-port or list of sub-ports.

```
sub-port show sub-port <sub-port-name-list> statistics
```

where

<sub-port-name-list>	is the particular sub-port or list of sub-ports that you want to display statistics for.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example shows sample output for the sub-port show parent-port statistics command.

```
> sub-port show statistics

+--------------------------- SUB-PORT STATISTICS SUMMARY --------------------+
| Sub-Port      | Byte                                | Pkt                        |
|               | Tx              | Rx                | Tx                | Rx         |
+---------------+-----------------+-----------------+-----------------+----------+
| sp1.v101.p9   | 574563251000    | 0                 | 574563253       | 0        |
| sp2.x102.p9   | 0               | 574563260000      | 0               | 574563260|
| sp3.v103.p10  | 574563271000    | 0                 | 574563273       | 0        |
+---------------+-----------------+-----------------+-----------------+----------+
```

This example shows sample output for the sub-port show parent-port <parent-port-name> statistics command.

```
> sub-port show parent-port 9 statistics

+------------------------- SUB-PORT STATISTICS FOR PARENT 9------------------------+
| Sub-Port      | Byte                                | Pkt                                |
|               | Tx              | Rx                | Tx                | Rx             |
+---------------+-----------------+-----------------+-----------------+----------------+
| sp1.v101.p9   | 574563251000    | 0                 | 574563253       | 0              |
| sp2.x102.p9   | 0               | 574563260000      | 0               | 574563260      |
+---------------+-----------------+-----------------+-----------------+----------------+
```

This example shows sample out put for statistics display command on a single sub-port.

```
sub-port show sub-port sp1.v101.p9,sp2.x102.p9 statistics
+----------- SUB-PORT sp1.v101.p9 STATISICS -----------+
| Direction     | Bytes               | Packets           |
+---------------+--------------------+--------------------+
| Rx            | 0                  | 0                  |
| Tx            | 4604013142000      | 4604013152         |
+---------------+--------------------+--------------------+
+----------- SUB-PORT sp2.x102.p9 STATISICS -----------+
| Direction     | Bytes               | Packets           |
+---------------+--------------------+--------------------+
| Rx            | 4604013198000      | 4604013201         |
| Tx            | 0                  | 0                  |
+---------------+--------------------+--------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 7-41
Clearing statistics for sub-port interfaces

You can clear statistic counts for all sub-port interfaces on a device or for a specific list of sub-port interfaces.

Step	Action
1	Clear statistics counts for sub-port interfaces. `sub-port clear sub-port <sub-port-object-list> statistics` where <sub-port-object-list> is a list of sub-ports which have already been configured on the device.

—end—

Example

This example clears statistics counts for a single sub-port interface:

```
sub-port clear sub-port sp1.v101.p9
```


Copyright© 2022 Ciena® Corporation

Procedure 7-42
Configuring a virtual switch for remote management and customer data traffic transiting a device

Configure a VLAN-emulating virtual switch for remote management traffic and virtual switches for customer data traffic to be encapsulated into a Q-in-Q network.

Step	Action
1	Set the virtual switch transform mode for all UNI ports to i-stamp:push,e-match-pop:stamp: `port set port <portNameList> vs l2-transform i-stamp:push,e-match-pop-stamp` where port <PortNameList> — is the name of the ports.
2	If required, add reserved VLANs for management and customer virtual switches: `virtual-switch add reserved-vlan <VLANIDList>` where reserved-vlan <VLANIDList> — is the name of the VLANs.
3	If the UNI port is currently a member of the remote management VLAN, remove it as a VLAN member: `vlan remove vlan <VLANIDList> port <PortNameList>` where vlan <VlanIDList> — is the name of the remote management VLAN. port <PortNameList> — is a list of the names of the UNI ports.
4	Create the virtual circuit using the remote management VLAN: `virtual-circuit ethernet create vc <vc> vlan <VlanIDList>` where vc <vc> — is the name of the virtual circuit. vlan <VlanIDList> — is the name of the remote management VLAN.

Copyright© 2022 Ciena® Corporation

5 Create the VLAN(s) to be used as the Q-in-Q SVLAN/EVC for the customer data traffic:

```
vlan create vlan <VlanIDList>
```

where

vlan <VlanIDList> is the name of the VLAN(s).

6 Add the NNI port(s) to the customer data traffic VLAN(s):

```
vlan add vlan <VLANIDList> port <PortNameList>
```

where

vlan <VlanIDList> is the name of the VLAN(s).

port <PortNameList> is the name of the port(s).

7 Create the virtual circuit(s) for each customer data EVC:

```
virtual-circuit ethernet create vc <vc> vlan <VlanIDList>
```

where

vc <vc> is the name of the virtual circuit.

vlan <VlanIDList> is the name of the VLAN used as the Q-in-Q SVLAN/EVC.

Note: Repeat this step for customer data EVC.

8 Create the virtual switches for management traffic and all customer data EVCs:

```
virtual-switch ethernet create vs <vs> vc <vc>
```

where

vs <vs> is the name of the virtual switch.

vc <vc> is the name of the virtual circuit.

Note: Repeat this step for the remote management VS and each customer data EVC.

9 Disable the LLDP on the UNI port.

```
lldp port <port number> disable
```


Copyright© 2022 Ciena® Corporation

10 Create the EVPL virtual switch members with translate-tag 0 for the management traffic:

```
virtual-switch ethernet add vs <VirtualSwitchName> port
<PortNameList> vlan <VlanIDList> translate-tag 0
```

where

vs <VirtualSwitchName	is the name of the virtual switch.
port <PortNameList>	is the name of the UNI port.
vlan <VlanIDList>	is the name of the remote management VLAN.

Note: Repeat this step for each UNI port to contain remote management traffic.

11 Create the virtual switch members for the customer data traffic (EPL or EVPL). For EVPL, complete step 12. For EPL complete this step:

```
virtual-switch ethernet add vs <VirtualSwitchName> port
<PortNameList>
```

where

vs <VirtualSwitchName>	is the name of the virtual switch.
port <PortNameList>	is the name of the UNI port.

12 Create the virtual switch members for the customer data traffic EVPL:

```
virtual-switch ethernet add vs <VirtualSwitchName> port
<PortNameList> vlan <VlanIDList> translate-tag
<TranslateTag>
```

where

vs <VirtualSwitchName	is the name of the virtual switch.
port <PortNameList>	is the name of the UNI port.
vlan <VlanIDList>	is the VID of the customer traffic received and transmitted on the UNI port.
translate-tag <TranslateTag>	is the VID of the customer traffic in the C-Tag when the frame is Q-in-Q encapsulated.

Note 1: For a non-translate action of the CVID, set the translate-tag attribute value to the same value as the vlan attribute.

Note 2: Repeat this step for each customer data EVC.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example shows remote management and customer traffic transiting the device through specific virtual switches and Q-in-Q EVCs.

Figure 7-41
Remote management and customer traffic example

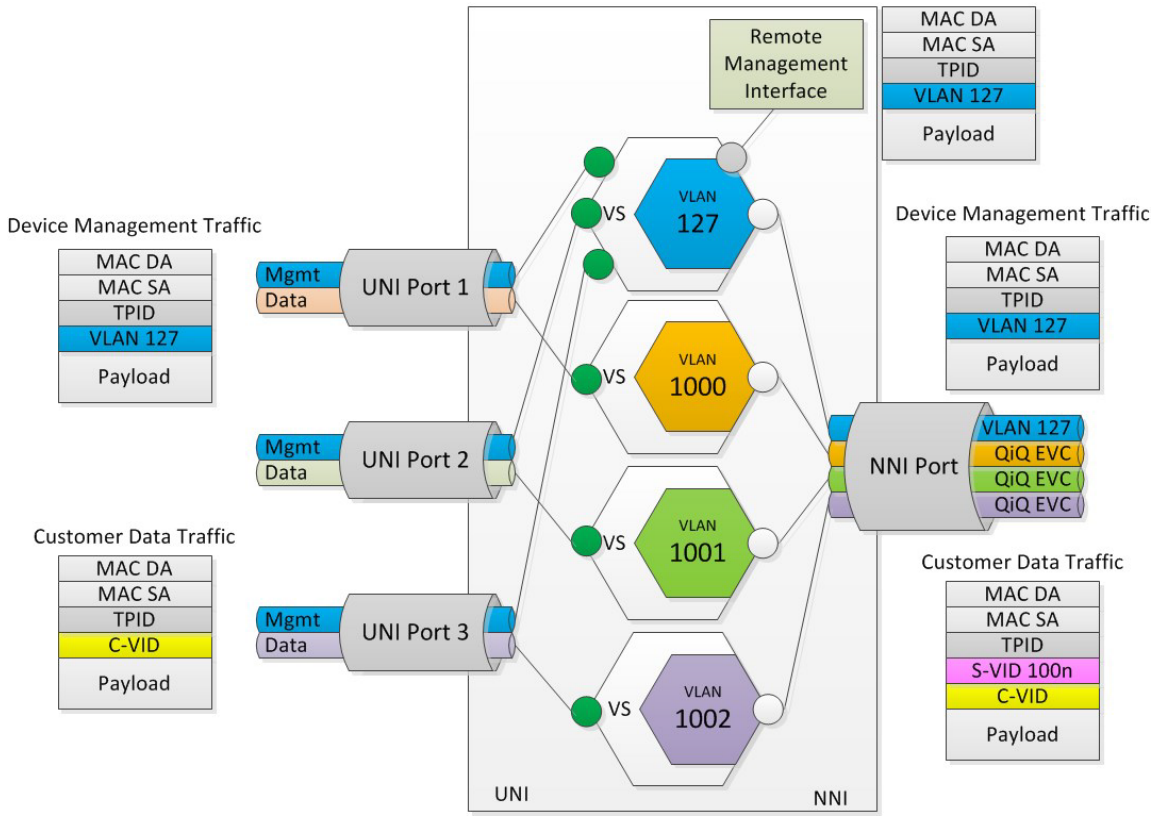

Copyright© 2022 Ciena® Corporation

This example configuration shows what is displayed in the Remote management and customer traffic example figure. Each of the virtual switch members for customer data traffic received on the three UNI ports are all EPL virtual switch members. In this example, the management VLAN exists as VLAN 127 and has all NNI ports as members.

```
port set port 1-3 vs-l2-transform i-stamp:push,e-match-
pop:stamp
vlan remove vlan 127 port 1-3
virtual-circuit ethernet create vc vc127 vlan 127
vlan create vlan 1000-1002
vlan add vlan 1000-1002 port 24
virtual-circuit ethernet create vc vc1000 vlan 1000
virtual-circuit ethernet create vc vc1001 vlan 1001
virtual-circuit ethernet create vc vc1002 vlan 1002
virtual-switch ethernet create vs vs127 vc vc127
virtual-switch ethernet create vs vs1000 vc vc1000
virtual-switch ethernet create vs vs1001 vc vc1001
virtual-switch ethernet create vs vs1002 vc vc1002
lldp set port 1 disable
virtual-switch ethernet add vs vs127 port 1 vlan 127 translate-
tag 0
lldp set port 2 disable
virtual-switch ethernet add vs vs127 port 2 vlan 127
translate-tag 0
lldp set port 3 disable
virtual-switch ethernet add vs vs127 port 3 vlan 127
translate-tag 0
virtual-switch ethernet add vs vs1000 port 1
virtual-switch ethernet add vs vs1001 port 2
virtual-switch ethernet add vs vs1002 port 3
```


Copyright© 2022 Ciena® Corporation

Procedure 7-43
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 port <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 port <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	Create a static Ethernet virtual circuit for the specified VLAN object: `virtual-circuit ethernet create {vc <name>} {vlan <Vlan>}[l2-transform <vc-transform-name>]`
6	Create a virtual switch instance and associate a virtual circuit: `virtual-switch ethernet create vs <vs> [vc <vc>]`
7	Disable the LLDP on the UNI port. `lldp set port <port number> disable`
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>] [pcp <NUMBER: 0-7>]` **—end—**

Copyright© 2022 Ciena® Corporation

Example

The example illustrated in the VS L2 transform: i-push,e-pop for Q-in-Q figure 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 7-42
VS L2 transform: i-push,e-pop for Q-in-Q

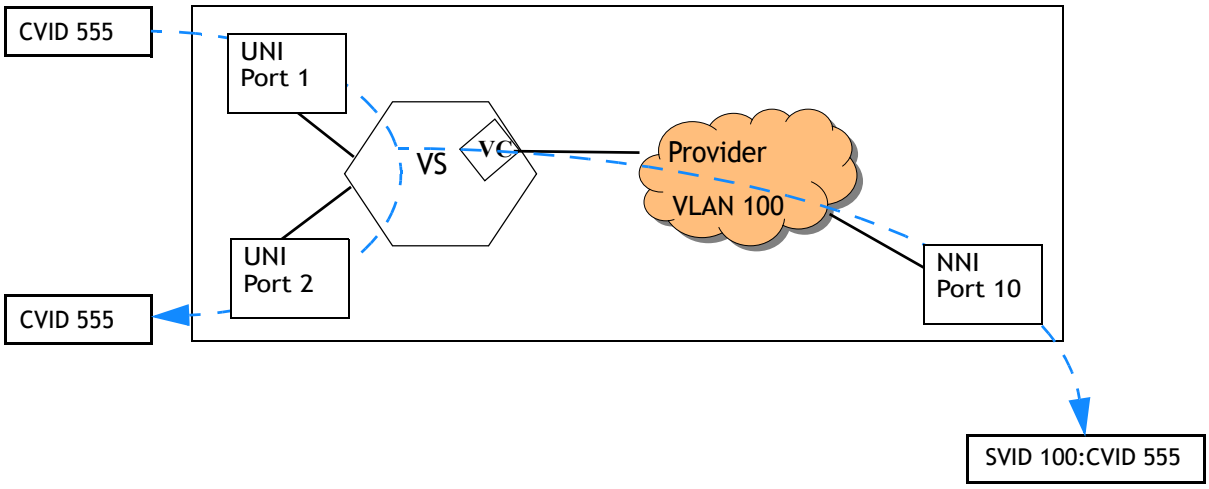

This example creates the configuration.

```
> 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-circuit ethernet create vc vc100 vlan 100
> virtual-switch ethernet create vs vs100 vc vc100
> lldp set port 1 disable
> lldp set port 2 disable
> virtual-switch ethernet add vs vs100 port 1 vlan 555
```


Copyright© 2022 Ciena® Corporation

Procedure 7-44
Configuring MPLS encapsulation of MEF8 TDM PWs

The restrictions for the SFP TDM ports are as follows:

- Native MPLS and MPLS encapsulation cannot be applied to the same SFP TDM port.
- Once a VLAN is used for MPLS encapsulation for the SFP TDM port, the VLAN cannot be used for MEF8 for that SFP TDM port.
- VLAN can only be used for one MPLS VS that is attached providing that all the tdm vcs are associated to the same TDM port. The tdm vcs can be attached to the same VPWS MPLS VS.
- VPWS MPLS VSs can have TDM VC bundles attached as long as all the TDM VCs are associated to the same TDM port. The TDM VCs can be attached to the same VPWS MPLS VS.
- The tdm-vc can only be attached to the MPLS VS or a VS cross-connect but not both.
- Ethernet ports cannot be attached to an MPLS VS if a tdm-vc is already attached and vice versa.

The restrictions for the PWE module or enhanced module TDM ports is similar to the SFP ports except that it applies across all PWE module or enhanced module TDM ports:

- Native MPLS and MPLS encapsulation are mutually exclusive in the ports of the PWE module or enhanced module port.
- Once a VLAN is specified for MPLS encapsulation, it cannot be used for MEF 8 for other PWE module or enhanced module TDM ports.
- VLAN can only be used for one MPLS VS in all PWE module or enhanced module TDM ports.
- The tdm-vc can only be attached to an MPLS VS or a VS cross-connect, but not both.
- For VPWS MPLS VSs, providing all the tdm vcs are associated with PWE module or enhanced module TDM ports, then the tdm vcs can be attached to the same VPWS MPLS VS.
- Ethernet ports cannot be attached to an MPLS VS if a tdm-vc is already attached and vice versa.

Copyright© 2022 Ciena® Corporation

Step	Action
1	Create the MPLS virtual switch: `virtual-switch ethernet create vs <vsName> mode <vpws|vpls>` where vs <vsName> — is the virtual switch name. mode <vpws|vpls> — sets the capability of the virtual switch.
2	Add the TDM virtual circuit to the MPLS virtual switch with a specified VLAN: `virtual-switch ethernet add vs <vsName> tdm-vc <tdm-vcName> vlan <1..4095>` where vs <vsName> — is the virtual switch name. tdm-vc <tdm-vcName — is the name of the TDM virtual circuit. vlan <1..4095> — identifies which packets to accept into the virtual switch.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 7-45
Configuring the CVID range as a single EVPL

Configure the CVID range as a single EVPL using the vlan-range-type option.

Note: The vlan-range-type is an optional parameter that can only be used when adding a new member. Set/unset operations are not applicable.

Step	Action
1	Configure the CVID range type as a single EVPL: `virtual-switch ethernet add vs <vs> port <port> vlan <VLAN ID list> vlan-range-type <strict|group>` where vs <vs> — is the virtual switch name. port <port> — is the port. vlan <vlan name> — is the VLAN name/ID. vlan-range-type <vlan-range type> — is the type specified for the vlan-range. **—end—**

Copyright© 2022 Ciena® Corporation

Example

This example configures the complete supported vlan IDs of 1-4094 with the vlan-range type as strict.

```
virtual-switch ethernet add vs vs1 port 2 vlan 1-4094 vlan-range-type strict
```

The user can configure the complete supported range as illustrated above or configure individual ranges as shown below.

```
virtual-switch ethernet add vs vs1 port 2 vlan 1-499 vlan-range-type strict
virtual-switch ethernet add vs vs1 port 2 vlan 500-999 vlan-range-type strict
virtual-switch ethernet add vs vs1 port 2 vlan 4000-4094 vlan-range-type
strict
virtual-switch ethernet add vs vs1 port 3 vlan 1-4094 vlan-range-type strict
virtual-switch ethernet add vs vs1 port 2 vlan 1000-1002 vlan-range-type
strict
```

Even though the last configuration is valid, it will consume one of the eight range resources available for port 2, which might not be considered an optimal usage of range resources.

This example configures an EVPL with the vlan-range-type as group.

```
virtual-switch ethernet add vs vs1 port 2 vlan 1-1000 vlan-range-type group
```


Copyright© 2022 Ciena® Corporation

Procedure 7-46
Setting the VLAN range for a port

Set the VLAN range for a port.

Step	Action
1	Set the vlan-range for a specific port: `port set port <port> vlan-range-index <NUMBER: 1..8> vlan-range <VLAN ID List>` **—end—**

Example

This example sets the first vlan-range on port 2.

```
port set port 2 vlan-range-index 1 vlan-range 100-199
```


Copyright© 2022 Ciena® Corporation

Procedure 7-47
Unsetting the VLAN range for a port

Unset the VLAN range for a port.

Step	Action
1	Unset the vlan-range for a specific port: `port unset port <port> vlan-range-index <NUMBER: 1..8>` **—end—**

Example

This example unsets the first vlan-range on port 2.

```
port unset port 2 vlan-range-index 1
```


Copyright© 2022 Ciena® Corporation

Procedure 7-48
Displaying the VLAN range for a device

Display the VLAN range for a device as required for the network plan.

Step	Action

1	Display the vlan-ranges used by a port/ports:
	`port show vlan-range-indexes`
	`port show port <port> vlan-range-indexes`
	—end—

Example

This example shows the vlan-range configured on port 2.

```
port show port 2 vlan-range-indexes

+------------------- PORT 2 VLAN RANGE INFO --------------------+
| Index          | Admin                  | Oper                   |
+----------------+------------------------+------------------------+
| 1              | 100-199                | 100-199                |
| 2              |                        |                        |
| 3              |                        |                        |
| 4              |                        |                        |
| 5              |                        |                        |
| 6              |                        |                        |
| 7              |                        |                        |
| 8              |                        |                        |
+----------------+------------------------+------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 7-49
Configuring a Q-in-Q virtual switch with SVLAN transform

Configure a Q-in-Q virtual switch with SVLAN transform by creating a virtual circuit with the vlan-transform option i-push,e-pop.

Note 1: Only applies to 390x, 392x, 3942, 5142, and 5160.

Note 2: Set/unset operations are not applicable.

Step	Action
1	Create an SVLAN: `vlan create vlan <vlan>`
2	Add the NNI port in the SVLAN: `vlan add vlan 100 port <port>`
3	Create a virtual circuit with vlan-transform: `virtual-circuit ethernet create vc <vc> vlan <vlan> vlan-transform <i-push,e-pop>` where vlan <vlan> — is the VLAN name. port <port> — is the port. vc <vc> — is the virtual circuit name. vlan-transform <i-push,e-pop> — is the SVLAN transform.
4	Create a virtual switch: `virtual-switch ethernet create vs <vs> vc <vc>`
5	Add the port as an EVPL member in the virtual switch: `virtual-switch ethernet add vs <vs> port <port> vlan <vlan>`
6	Add the port as an EVPL member with a VID range in the virtual switch: `virtual-switch ethernet add vs <vs> port <port> vlan <range> vlan-range-type <strict|group>`

—end—

Copyright© 2022 Ciena® Corporation

Example: Configuration with vlan-range set only on the UNI port

This example create the following configuration:

SVLAN:100, UNI Port:1, NNI Port: 2

cVids: 10, 20 and range 100…199

A range resource is used at the UNI port.

```
vlan create vlan 100
vlan add vlan 100 port 2
virtual-circuit eth create vc transparen_vc vlan 100 vlan-transform i-push,e-
pop
virtual-switch eth create vs transparent_vs vc transparen_vc
virtual-switch ethernet add vs transparent_vs port 1 vlan 10
virtual-switch ethernet add vs transparent_vs port 1 vlan 20
virtual-switch ethernet add vs transparent_vs port 1 vlan 100-199 vlan-range-
type strict
```

Example: Configuration with vlan-range set at the NNI port

This example optimizes H/W resources by using a range resource at the NNI port also. This can help in optimizing overall H/W resource usage.

```
vlan create vlan 100
vlan add vlan 100 port 2
virtual-circuit eth create vc transparent_vc vlan 100 vlan-transform i-push,e-
pop
virtual-switch eth create vs transparent_vs vc transparent_vc
virtual-switch ethernet add vs transparent_vs port 1 vlan 10
virtual-switch ethernet add vs transparent_vs port 1 vlan 20
port set port 2 vlan-range-index 1 vlan-range 100-199
virtual-switch ethernet add vs transparent_vs port 1 vlan 100-199 vlan-range-
type strict
```

Example: Configuration without vlan-translation set to i-push, e-pop

In this example, all eight range resources of UNI port 1 are used in following configuration example. Since the virtual-circuit is not created with vlan-translation set to i-push,e-pop, the device does not consume any range resources for NNI port.

```
vlan create vlan 100
vlan add vlan 100 port 2
virtual-circuit eth create vc vc1 vlan 100
virtual-switch eth create vs vs1 vc vc1
virtual-switch ethernet add vs vs1 port 1 vlan 10
virtual-switch ethernet add vs vs1 port 1 vlan 20
virtual-switch ethernet add vs vs1 port 1 vlan 100-199 vlan-range-type strict
virtual-switch ethernet add vs vs1 port 1 vlan 200-299 vlan-range-type strict
virtual-switch ethernet add vs vs1 port 1 vlan 300-399 vlan-range-type strict
virtual-switch ethernet add vs vs1 port 1 vlan 400-499 vlan-range-type strict
virtual-switch ethernet add vs vs1 port 1 vlan 500-599 vlan-range-type strict
virtual-switch ethernet add vs vs1 port 1 vlan 600-699 vlan-range-type strict
virtual-switch ethernet add vs vs1 port 1 vlan 700-799 vlan-range-type strict
virtual-switch ethernet add vs vs1 port 1 vlan 800-899 vlan-range-type strict
```


Copyright© 2022 Ciena® Corporation

Example: Configuration with VLAN range set to group

In the example below, the CVID range is configured on same UNI port, and it consumes individual H/W resources for each CVID present in range. Since the vlan-range option is set to group, the same VLAN range can be configured on NNI port to save H/W resources for the NNI port.

The device uses all eight range resources of NNI port also. This reduces overall H/W resource usage.

```
vlan create vlan 1000
vlan add vlan 1000 port 2
virtual-circuit eth create vc transparent_vc vlan 100 vlan-transform i-push,e-
pop
virtual-switch eth create vs vs2 vc transparent_vc
port set port 2 vlan-range-index 1 vlan-range 1000-1099
virtual-switch ethernet add vs vs2 port 1 vlan 1000-1099 vlan-range-type group
port set port 2 vlan-range-index 2 vlan-range 1100-1199
virtual-switch ethernet add vs vs2 port 1 vlan 1100-1199 vlan-range-type group
port set port 2 vlan-range-index 3 vlan-range 1200-1299
virtual-switch ethernet add vs vs2 port 1 vlan 1200-1299 vlan-range-type group
port set port 2 vlan-range-index 4 vlan-range 1300-1399
virtual-switch ethernet add vs vs2 port 1 vlan 1300-1399 vlan-range-type group
port set port 2 vlan-range-index 5 vlan-range 1400-1499
virtual-switch ethernet add vs vs2 port 1 vlan 1400-1499 vlan-range-type group
port set port 2 vlan-range-index 6 vlan-range 1500-1599
virtual-switch ethernet add vs vs2 port 1 vlan 1500-1599 vlan-range-type group
port set port 2 vlan-range-index 7 vlan-range 1600-1699
virtual-switch ethernet add vs vs2 port 1 vlan 1600-1699 vlan-range-type group
port set port 2 vlan-range-index 8 vlan-range 1700-1799
virtual-switch ethernet add vs vs2 port 1 vlan 1700-1799 vlan-range-type group
```


Copyright© 2022 Ciena® Corporation

Procedure 7-50
Displaying information related to Q-in-Q virtual switch with SVLAN transform

Display a virtual circuit with vlan-transform.

Step	Action
1	Display a specific vlan-transform: `virtual-circuit ethernet show vc <vc-name>`

—end—

Example

This example shows the output from the virtual-circuit ethernet show command for virtual circuit 2000.

```
virtual-circuit ethernet show vc vc#2000

+-------------- ETHERNET VIRTUAL CIRCUIT vc#2000 INFO ----------------+
| Parameter                | Value                                   |
+--------------------------+-----------------------------------------+
| Name                     | vc#2000                                 |
| ID                       | 1                                       |
| Provider VLAN            | (2000)  VLAN#2000                       |
| L2 Transform             |                                         |
| Provider VLAN Transform  | i-push,e-pop                            |
| Psuedo-Wire Type         | untagged                                |
| VS Using                 |                                         |
| Stats Collection         | Off                                     |
+--------------------------+-----------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Quality of Service configuration

This section explains how to configure Quality of Service (QoS) features of the system software. Topics are:

- “Class of Services (CoS) policies and mapping” on page 8-1
- “RCoS and RColor behavior for QinQ virtual switch” on page 8-4
- “Traffic profiling” on page 8-10
- “Traffic profiling with VLAN range support” on page 8-20
- “Traffic profiling with hierarchical ingress metering” on page 8-21
- “MEF 10.3 ingress metering” on page 8-27
- “Congestion management” on page 8-48
- “Egress queuing and scheduling” on page 8-50
- “Local CPU-generated control, management and OAM traffic QoS treatment” on page 8-61
- “Dynamic microflow detection and greedy flow monitoring” on page 8-73

Class of Services (CoS) policies and mapping

The system software associates an internal resolved CoS (RCoS) value and resolved Color (RColor) with every ingress frame that classifies to a port. The RCoS and RColor values provide the baseline for CoS treatment of a frame as it switches through the device. The RCoS value ranges from 0-7, where 7 has the highest priority and 0 has the lowest priority. RColor values are green and yellow. The initial RCoS and RColor values depend on the configuration of the of the port, including virtual switch memberships.

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.

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 RCoS and RColor is determined by the encap CoS policy (and priority)

Copyright© 2022 Ciena® Corporation

that is configured at the virtual switch. When the encap CoS policy configured on a virtual switch member is port-inherit, the resolved Cos Policy of the port is applied. The system does not add a port as a virtual switch member if the port resolved CoS policy is configured based on the inner VLAN tag PCP and DEI bits.

The 3924, 3926 and 3928 devices do not allow the mixed encap-cos-policy configuration. The 3942, 5142 and 5160 devices allow this configuration, but display a warning message when one of these configurations is attempted:

- One membership with encap-cos-policy port inherit and another membership with VS-inherit
- One membership with port-inherit encap-cos-policy and another membership with fixed encap-cos-policy

All VS memberships (EPL/EVPL) of a port have either port-inherit or fixed/vs-inherit encap-cos-policy.

> ***Note:*** If an EVPL member has its port-inherit policy set to l3-dscp, then all frames containing an IP header in the other EVPLs configured on the same client port follow the L3 mapping table, even if set to vs-inherit or fixed. If the frames are not IP and are tagged, then the L2 portion of the mapping table is in effect. If the frames are untagged without an IP header, then the RCoS and RColor 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 RCoS and RColor values.

It shows which policy and portion of the map is in effect depending on the port's configuration.

Copyright© 2022 Ciena® Corporation

Figure 8-1
CoS policy and mapping

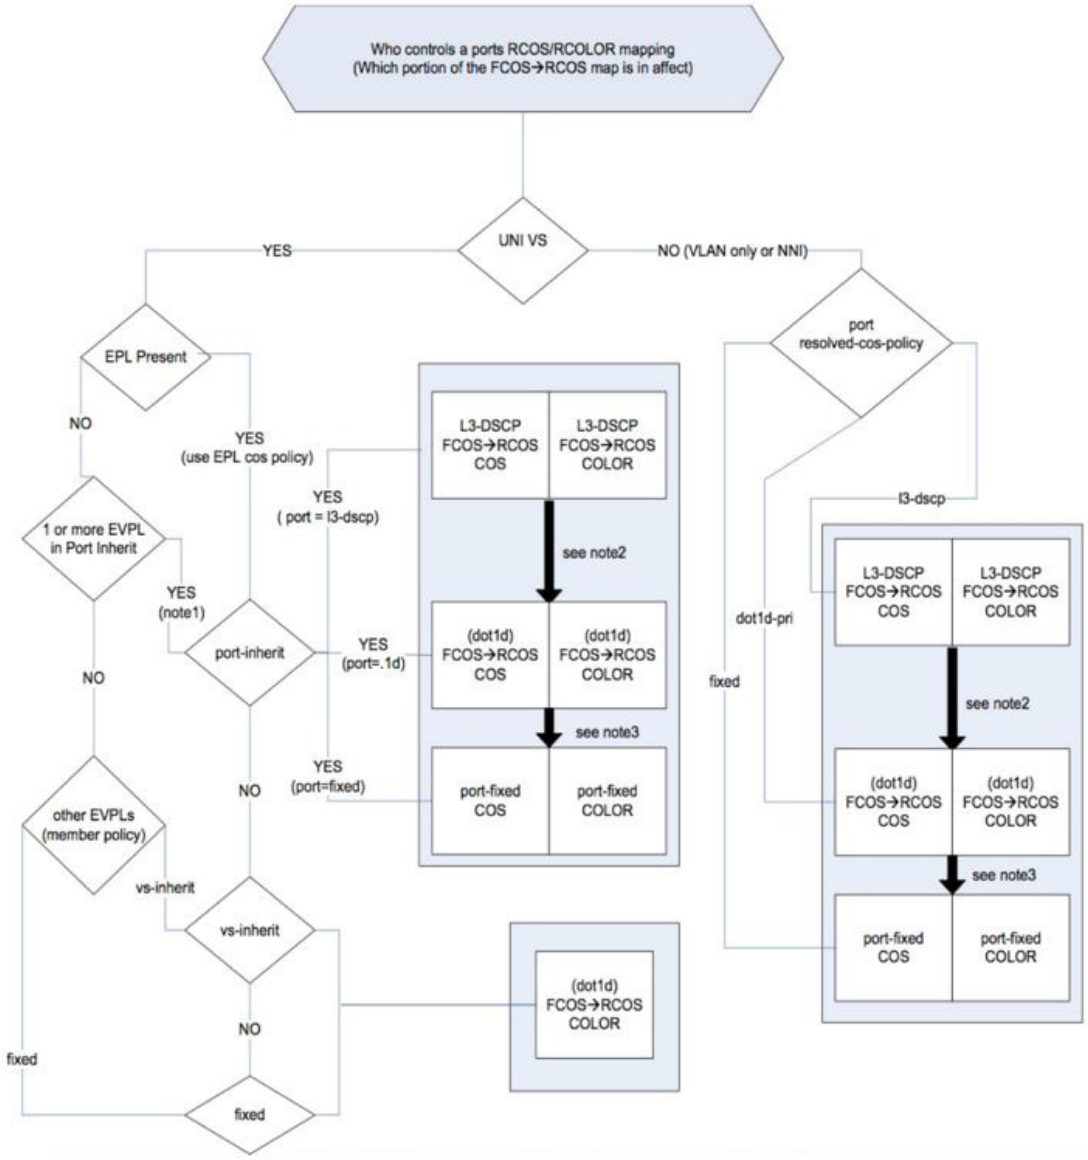

Note 1: An EVPL in port inherit influences all EVPL entries behaviors. If an EVPL member has its port-inherit policy set to l3-dscp, then all frames containing an IP header in the other EVPLs configured on the same client port follow the L3 mapping table, even if set to vs-inherit or fixed.

Note 2: If the EVPL is not a layer 3 frame and tagged, use the L2 portion of the table. If not tagged and no IP header in the frame, the cos and color are fixed to the port's fixed settings.

Copyright© 2022 Ciena® Corporation

Note 3: If the EVPL is an untagged (non-confirming) frame, the CoS and color is fixed with the port's fixed settings.

RCoS and RColor behavior for QinQ virtual switch

The QinQ virtual switch: Untagged traffic table describes the behavior for untagged traffic RCoS and RColor behavior for QinQ virtual switch. This example captures the untagged frame's RCoS/RColor and outer S-tag PCP/DEI.

Table 8-1
QinQ virtual switch: Untagged traffic

VS member attributes			VS fixed CoS	Ingress port RCoS policy	Ingress port fixed RCoS	Ingress port fixed RColor	Results		
EP/ EVPL	Encap CoS policy	Fixed CoS					RCos/S-PCP	RColor	S-DEI
EPL	Fixed	4	2	no effect	6	Yellow	4 (VS memb fixed CoS)	Yellow (Ingress port fixed RColor)	0
EPL	VS-inherit	4	2	no effect	6	Yellow	2 (VS Fixed CoS)	Yellow Ingress port fixed RColor)	0
EPL	Port-inherit	4	2	no effect	6	Yellow	6 (Ingress port fixed RCoS)	Yellow (Ingress Port Fixed RColor)	0
EVPL	no effect	no effect	2	no effect	6	Yellow	2 (VS fixed CoS)	Yellow (Ingress Port Fixed RColor)	0

Copyright© 2022 Ciena® Corporation

This table describes the behavior for QinQ RCoS and RColor behavior for the 3924, 3926 and 3928 switches.

Table 8-2
QinQ RCoS and RColor behavior for 3924, 3926 and 3928

Incoming frame type	Platform	VS member attributes		Ingress port RCoS policy	Results			
		EP/ EVPL	Encap CoS policy		S-PCP	S-DEI	RCoS	R-Color
Tagged (C=PCP= 3 C-DEI=1)	3924, 3926, 3928	EVPL/ EPL	Fixed (=4)	no effect	Fixed (=4)	0	Fixed (=4)	Green
Tagged (C-PCP=3 C-DEI=1)		EVPL/ EPL	VS-inherit (VS fixed CoS=2)	no effect	VS Fixed CoS (=2)	0	VS fixed CoS (=2)	Green
Tagged (C-PCP=3 C-DEI=1		EVPL/ EPL	Port-inherit	dot1d map	C-PCP (=3)	0	dot1d Map (Key# C-PCP, C-DEI)	dot1d Map (Key# C-PCP, C-DEI)
Tagged (C-PCP=3 C-DEI=1)		EVPL/ EPL	Port-inherit	Fixed	C-PCP (=3)	0	Ingress port fixed RCoS	Ingress port fixed RColor

Ports with virtual switch members

If an untagged-data-vs is assigned to the port, then the policy is based on the sum of all of the VS member encap-cos-policy settings that are present on that port. If any of the EPL VS members on the port have an encap-cos-policy of port-inherit, then that policy is selected and the assignment of RCoS is the same as for the ports that do not have VS members. In this case, any EVPLs with any other encap-cos-policy settings are ignored. If the VS member is an EPL, then the RCoS is determined by the encap-cos-policy of that member.

Copyright© 2022 Ciena® Corporation

This table describes the encapsulation cos policy with respect to resolved CoS:

Table 8-3
Encap-cos-policy

EPL resolved-cos-policy	RCoS
fixed	Use encap-fixed-dot1 dpri of VS member
port-inherit	Use encap-fixed-dot1 dpri of VS member
vs-inherit	Use encap-fixed dot1dpri of VS

If all the VS members on the port are EVPLs, the RCoS is set to the encap-fixed-dot1 dpri of the VS. If the encap-cos-policy of one of the EVPLs is set to port-inherit, and the resolved cos policy of the port is l3-dscp-cos and the untagged frame has a valid DSCP priority, then the RCoS -> RCoS map is used.

Remarking behavior for VLAN switching

This table describes the remarking behavior for VLAN switching for 3924, 3926, 3928, 3942, 5142 and 5160 devices.

Table 8-4,
Remarking behavior: 3924, 3926, 3928, 3942, 5142, 5160

Platform	Ingress port: 'resolved-cos-remark-L2'	Egress port: 'egress frame cos policy'	Egress queue	Egress outer-tag P-bit	Egress outer-tag DEI bit
3924, 3926 3928 3942 5142 5160	false	ignore	The packet is put into the proper queue based on FCos -> RCos table resulting in RCoS -> queue table mapping both being on the ingress port.	No remarking	No remarking
		rcos-to-l2-outer-pcp-map		No remarking	No remarking
	true	ignore		No remarking	No remarking
		rcos-to-l2-outer-pcp-map		Remarked based on FCoS -> RCos then RCos -> FCos tables mapping	Remarked based on FCoS -> RCos then RCos -> FCos tables mapping
Note: No remarking means that the ingress outer tag p-bit or CFI/DEI bit is not changed after the packet leaving the box.					

Copyright© 2022 Ciena® Corporation

The resolved CoS Policy options are:

- Outer .1D mapped (dot1d-tag1-cos) -The Priority Code Point (PCP)/Layer 2 (L2) CoS 802.1D priority value from the outer-tag is mapped to RCoS and RColor 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.
- Inner .1D mapped (dot1d-tag2-cos) - Provides the capability to derive a resolved CoS and color based on inner VLAN tag and PCP and DEI bits if the incoming frame is double tagged. This tag provides the support derivation of the internal resolved CoS and color from the inner VLAN tag. This gives the flexibility to determine whether the user wants to derive internal RCoS and RColor from the incoming frame inner tag or outer tag PCP/DEI value. If the port resolved CoS policy is configured to dot1d-tag2-cos, the incoming frame's internal RCos/RColor is derived from the inner tag. There is no behavior change on the internal RCos/RColor if the RCoS/RColor is derived from the traffic profiling.

  When a port is in resolved-cos-policy dot1d-tag2-cos mode, and a frame does not conform to .1D, the frame receives an RCoS equal to the "fixed resolve COS" set on the port. The RCOLOR of an untagged frame is the "fixed r-color value configured on the port. The RCOLOR of a single tagged or priority tagged frame is green.

  In the case scenario in the Use case figure, the CPE sends frames to the access network which then adds an outer VLAN tag, with VLAN ID X, as the frames move toward the service provider network. The service provider network then adds or modifies (translate tag) as an additional VLAN tag, with VLAN ID Z. The service provider performs QoS operations based on the original customer VLAN tag and not the tag that was added while traversing the access network. The access network may be the third-party network where the outer VLAN tag's PCP and DEI bit value is not necessarily mapped one-to-one with the frame coming from the CPE device.

Copyright© 2022 Ciena® Corporation

Figure 8-2
Use case

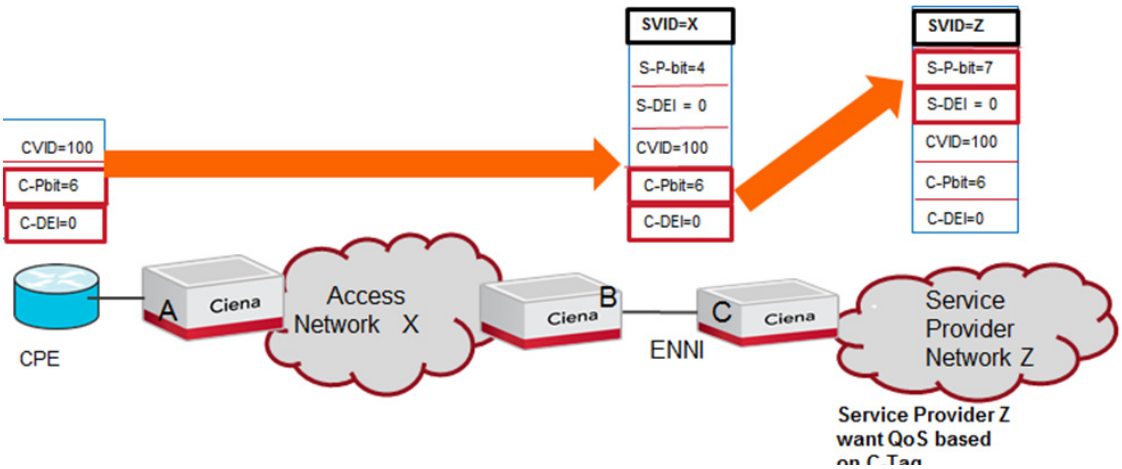

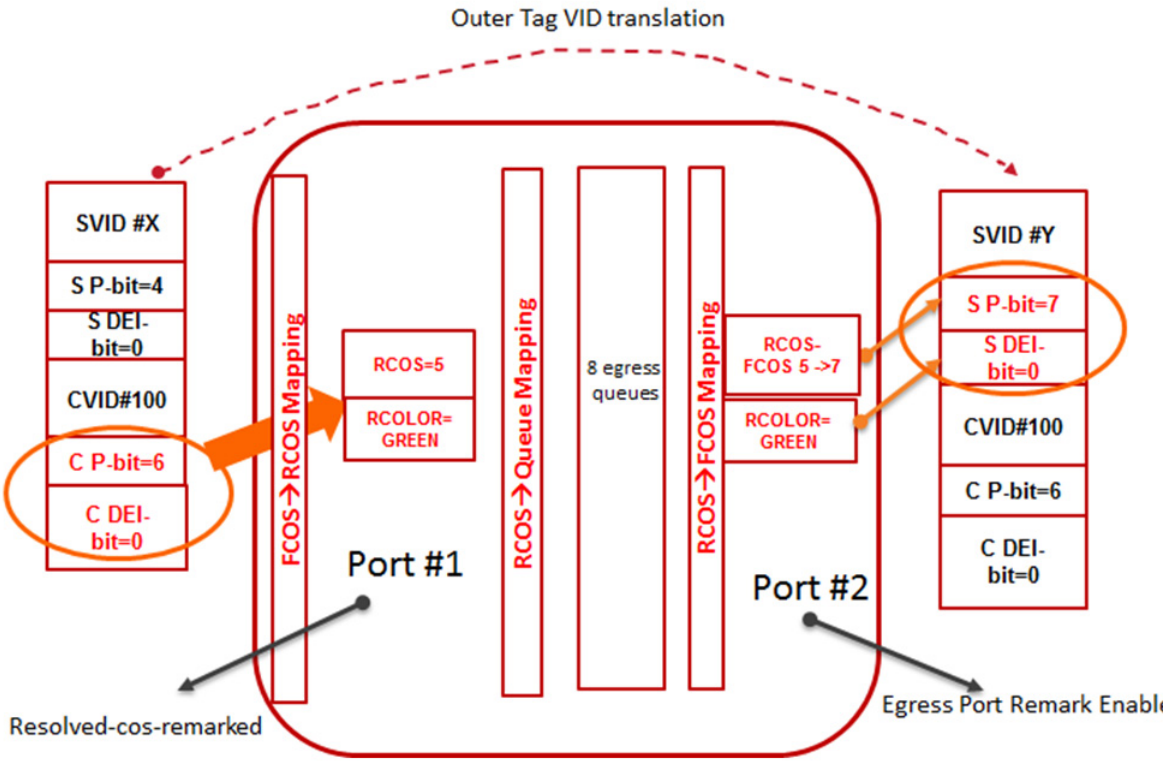

- Fixed CoS (fixed-cos), Fixed RCoS (fixed-rcos) and Fixed RColor (fixed-rcolor) - CoS values in the frame are ignored and fixed RCoS is applied to the frame from the Fixed Resolved CoS and Fixed Resolved Color value

Copyright© 2022 Ciena® Corporation

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 RCoS and RColor value from the Resolved CoS Mapping table. When this policy is applied, CoS mapping is in trusted mode.

RCoS to FCoS mappings for Ethernet networks

The number of RCoS -> FCoS mapping tables is equal to number of physical ports on devices with over 16 non-aggregate-ports and 16 for devices with less than or equal to 16 ports.

- For devices with less than or equal to 16 ports, SAOS supports 1 default, 3 “global” and 12 custom “l2-only” maps.
- For devices with more than 16 ports, SAOS supports 1 default, 3 “global”, and the remaining are custom “l2-only” maps.

This table describes the supported RCoS > FCoS mappings.

Table 8-5
Supported types

Type	Description
Global type (default)	Supports four maps (including default) for Ethernet networks and MPLS networks. The default “DefaultRcosFcos” is always global type. Cannot change from global type to L2-only type.
L2-only type	Supports 12 maps on devices with less than or equal to 16 ports. For devices with more than 16 ports, supports device number of ports -4 (1 default + 3 “global”). Cannot change from L2-only type to global type. L2-only maps are not attached to the port if that port is already part of an MPLS network. If an L2-only is already attached to a port, then that port is not added to the MPLS virtual switch. You cannot modify the mpls-tc field for “l2-only” map.

Copyright© 2022 Ciena® Corporation

Note: There is no difference between “global” type and “l2-only” type RCoS -> FCoS maps on non-MPLS supported devices.

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 on how the traffic flow compares to these rates, the R-COLOR of each frame is set to a specific value on ingress. Traffic ingressing at a rate:

- Up to CIR is marked Green and allowed through.
- Above CIR and less than PIR is marked Yellow and allowed through.
- Above PIR is marked Red and dropped.

Note: On a port, the CIR cannot be full port speed depending on Layer 1 (transport) or Layer 2 (payload). 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 can create a serious condition in which the device does not receive critical frames. To avoid this issue, the CIR must only be configured to a maximum of the port’s physical bandwidth - 1 Mbps.

Global profile mode

The global profile mode determines the type of profiles that can be created. There are three modes:

- global-l2-l3-extended
- global-l2-l3
- port

The default global profile mode is port. This mode allows you to configure standard profiles. Each standard profile is associated with a port. Standard profiles can also be configured to be hierarchical.

The global-l2-l3 mode allows traffic profiles to be configured on a global basis. These profiles do not require association to a port.

The global-l2-l3-extended mode is used for dynamic microflow detection on a global basis. See “Dynamic microflow detection and greedy flow monitoring” on page 8-73 for details. This mode adds support for extra classifiers over the global-l2-l3 mode.

Note the global profile mode is not related to the mode that can be set for the port under traffic profiling. That mode pertains to standard profiles when traffic profiling global mode is set to “port”. To switch between modes, no profiles can have been configured.

Copyright© 2022 Ciena® Corporation

Profile types

When the profile mode is set to global-l2-l3 or global-l2-l3-extended, an option to select between dynamic and global profile types is available. The profile type is set to “none” when the user sets the profile mode to port or global-l2-l3-extended. The profile type is set to “global” when the user sets the profile mode to global-l2-l3.

> ***Note:*** For SAOS 6.18, this option can only be used to set a “dynamic” profile type when the user sets the profile mode to global-l2-l3-extended.

Global traffic profiles

Global profiles cannot be configured to meter traffic from the CPU port and auxiliary ports. When a benchmark test is enabled, all global-profiles that apply to the port under test also apply to the benchmark port. On a benchmark disable, the affected global profiles applied on internal benchmark port are returned to their original state. Configuration changes to global profiles are not allowed if benchmark is enabled and the profile contains the port under test or the profile has no port classifier (all faceplate ports) and traffic profiling and the global profile are enabled. Benchmark supports VPWS, and hierarchy and classifier lists for global traffic profiles.

ACL deny rules take precedence over global traffic profiles. Even though the global profiles are set to classifier-mode “narrow”, packets that match an ACL deny rule count towards a matching global traffic profile meter as both are done in the IFP in parallel.

You can create, set and get global profiles by means of SNMP. You can also get associated statistics by means of SNMP.

> ***Note:*** All global profiles are deleted and the global profile mode is removed when downgrading to releases prior to SAOS 6.14.

Lists

List attributes let you have multiple classifier entries being metered by a common meter. The attributes of dot1dpri and dscp can be lists for a single global profile. For example, a list of dot1dpri of 1, 2 means that the meter matches on packets with dot1dpri of 1 or packets with dot1dpri of 2. The list restriction is the same as it is for standard traffic profiles. These attributes can be set to a list of values and a range of values. The VLAN attribute can be set to a list of four VLAN IDs.

> ***Note:*** The dot1dpri lists and dscp lists cannot be set in the same traffic profile. Because the 1 dscp and 1 dot1dpri value combination was allowed in previous releases for backwards compatibility, you can set 1 dscp value with a list of dot1dpri values. You can also set 1 dot1dpri value with a list of dscp values.

Copyright© 2022 Ciena® Corporation

The set and unset command operates on the dot1dpri, dscp and vlan lists the same way as for standard profiles. A set of values means an addition to the existing set of values. Any value that is in the set list that was already set is ignored. An unset of a list of values means the removal of the values that were listed from the classifier. Any value that was in the unset list that was not originally set is ignored.

The number of classifier resources used is 2 classifiers * number of dscp or dot1dpri values * number vlan values. If there are no vlan values, then it is 2 classifiers * number of dscp or dot1dpri values. If there are no lists, then the number of classifiers used is 2. For each classifier used, 2 counter resources are used.

> ***Note:*** Changing a classifier in a profile using the set or unset command may result in some packets not being accounted for in the statistics for that profile if the change is being done while the profile is enabled.

Hierarchy

Global profiles support a mix of shared BW hierarchical ingress metering (HIM) and non-shared BW metering profiles similar to standard profile HIM and advanced modes. The parent profiles are more flexible with global profiles. The parent profile supports these classifiers:

- Vlan
- Vs
- Port list
- Dot1dpri
- Dscp
- Base-etype
- Dst-ip
- Src-ip
- L4-dst-port
- L4-src-port
- Untagged

The parent profile only supports port lists. It does not support lists of dot1dpri, dscp or vlan. Any child profile of the parent cannot have classifiers that have been specified by the parent.

The parent profile meter can only set the cir and cbs. There is no pir/eir and ebs attribute that can be set as the parent profile meter is a single bucket meter.

Copyright© 2022 Ciena® Corporation

The child profile and the parent profile can be admin disabled or enabled. The child profile becomes enabled if the child profile, the parent profile and traffic profiling are all admin enabled. The child profile becomes disabled if any one of the child profile, parent profile or traffic profiling is admin disabled.

The parent profile does not use any classifier or counter resources. The parent profile uses two meter resources.

The child profile resource consumption is the same as a non-shared BW metering profile.

MPLS VS

Global profiles support VPLS, Ethernet QinQ VSes, and VPWS VSes. Note for VPWS VSes only, the UNI side of the VS is metered. This differs from VPLS and Ethernet QinQ VSes where both UNI and NNI side are metered.

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 on creation, the traffic profile does not work without configuring appropriate classifiers.

Port profiling mode

This table describes the classifiers for the port profiling mode.

Table 8-6
Classifiers

Classifier	Description
dot1dpri	Classifies to the traffic profile based on 802.1D priority value ranging from 0-7. Up to 8 classifiers are supported per port. Only one value can be specified for global profiles.
ip-prec	Classifies to the traffic profile based on IP precedence value ranging from 0-7. Up to 8 classifies are supported per port.
dscp	Classifies to the traffic profile based on DSCP. Up to 64 classifiers are supported per port. Both IPv4 and IPv6 are supported.
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. A VLAN list with more than one VLAN can be used with limited support. See "Traffic profiling with VLAN range support".

Copyright© 2022 Ciena® Corporation

Table 8-6
Classifiers

Classifier	Description
vs	Classifies the traffic profile based on the Ethernet VS ID. The global profile classifies on an MPLS virtual switch.
untagged	Classifies to the traffic profile if the traffic has no VLAN tag. Any other classifiers are ignored; they 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.
narrow	Drops packets that do not count towards the profile meter if there is a match. The meter mode for the global profile is always set to narrow.

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 on 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 this table.*

Table 8-7
Configurable classifiers

Classifier	Per mode	Per port (3903, 3903x, 3904, 3905, 3906)	Per port (3924, 3926, 3928, 3942)	Per port (5142, 5160)
dot1dpri	8	8	8	8
ip-prec (ipp)	8	8	8	8
dscp	64	64	64	64
vlan or vs*	1	1024	2048	4096
vlan or vs/dot1dpri*	8	1024	2048	4096

Copyright© 2022 Ciena® Corporation

Table 8-7
Configurable classifiers

Classifier	Per mode	Per port (3903, 3903x, 3904, 3905, 3906)	Per port (3924, 3926, 3928, 3942)	Per port (5142, 5160)
vlan or vs/ip-prec*	8	1024	2048	4096
vlan or vs/dscp*	64	1024	2048	4096
untagged	1	1	1	1

Note 1: The actual number of configurable classifiers is 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 a bout traffic profiling error codes, refer to "Traffic profiling error codes".

Note 3: *The VS maximum is up to the configurable limits of the system.

Global l2-l3 profiling mode

This table describes the classifiers for the global l2-l3 profiling mode.

Table 8-8
Classifiers

Classifier	Description
dot1dpri-list	Classifies to the traffic profile based on 802.1D priority value ranging from 0-7. Up to 8 classifiers are supported per port. Dotdpri lists are supported.
dscp-list	Classifies to the traffic profile based on DSCP. Up to 64 classifiers are supported per port. Both IPv4 and IPv6 are supported. Dscp lists are supported.
dscp-remark-policy	Leaves ingressing SVLAN 802.1p bit value as-is or marked with a fixed value.
vlan-list	Classifies to the traffic profile based on VLAN ID. A VLAN does not have to be created on the device to be part of a global profile. VLAN lists of up to four VLANs are supported. A VLAN list with more than four VLANs can be used with limited support. See "Traffic profiling with VLAN range support".

Copyright© 2022 Ciena® Corporation

Table 8-8
Classifiers

Classifier	Description
vs	Classifies the traffic profile based on the Ethernet VS ID. Note that traffic profiles cannot classify to an MPLS VS in port profiling mode. The global profile classifies on an MPLS virtual switch.
List of ports (includes aggregations)	This classifier can be configured on global l2-l3 profiles. An aggregation port cannot be deleted if it is associated with a global profile. If a physical port is part of a global profile, that port does not get removed from the global profile. The physical port is added to the classifier entry of all the global profiles that have the aggregation port specified in the profile. When a physical port is removed from the aggregation port, the physical port is removed from the classifier entry of all global profiles that have the aggregation port specified in the profile.
untagged	Classifies to the traffic profile if the traffic has no VLAN tag.
src-ip (IPv4) (CIDR notation)	Specifies the source IPv4 address to classify on.
dst-ip (IPv4) (CIDR notation)	Specifies the destination IPv4 address to classify on.
l4-src port	Is the layer 4 source port.
l4-dst-port	Is the layer 4 destination port.
Base-etype	Specifies the base Ethernet type to classify on. This classifier can be used to classify ARP packets.
IP-protocol	Is the IP protocol (0..255) that supports IPv4 and IPv6.
narrow	Drops packets that do not count towards the profile meter if there is a match. The meter mode for the global profile is always set to narrow.

Note 1: Other than the port list, all classifiers can only be a single value.

Note 2: There is no corresponding ACL filter for the untagged classifier.

Note 3: VPWS virtual switches are supported.

Copyright© 2022 Ciena® Corporation

Determining metering attributes

The meter provisioning mode can be set to either EIR or PIR. The mode that is chosen dictates how provisioning is accepted, displayed, and saved. If EIR is specified, PIR provisioning is disallowed. This table describes the metering attributes which apply to port and global l2-l3 profiles.

Note: The precedence attribute only applies to a global l2-l3 profile.

Table 8-9
Metering attributes

Metering attributes	Description
cir	Committed information rate in kbits ranging from 0 up to the administrative port speed. (Mandatory, if PIR is not specified.) Global parent profiles are single meters and can be provisioned with CIR and CBS only.
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 1:*** On 3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928, 3942, 5142, and 5160 switches, the configured CIR and PIR always normalize to a rate in varying increments depending on the sum of the configured rates of the CBS and EBS and the configured rate of the PIR. ***Note:*** SAOS 6.x devices do not support throughput beyond actual port speed. 6.x devices allow the user to configure a CIR value beyond the corresponding port speed. A warning message is displayed.

Copyright© 2022 Ciena® Corporation

Table 8-9
Metering attributes

Metering attributes	Description
	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 512 Kbps • 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 The burst rate is as follows: • 0 Kbytes <= cbs+ebs < 16M (16384 Kbytes, the burst/rate steps are 4K/64 kpbs • 16384 Kbytes <= burst-size < 32768 Kbytes, the burst/rate steps are 8K/128 Kbps burst/rate steps • 32768 Kbytes <= burst-size < 65536 Kbytes, the burst/rate steps are 16K/256 Kbps burst/rate steps • 65536 Kbytes <= burst-size < 131072 Kbytes, the burst/rate steps are 32K/512 Kbps burst/rate steps • 131072 Kbytes <= burst-size < 524288 Kbytes, the burst/rate steps are 64K/102.4 Kbps • 0 <= shaper-rare < 16.78 Gpbs, the burst/rate steps are 4K/64 Kbps burst/rate steps • 16.78 <= shaper-rate < 35.55 Gpbs, the burst/rate steps are 8K/128 Kbps burst/rate steps • 33.55 Gbps <= shaper=rate < 80 Gpbs, the burst/rate steps are 16K/256 Kbps burst/rate steps

Copyright© 2022 Ciena® Corporation

Table 8-9
Metering attributes

Metering attributes	Description
cbs	Committed Burst Size (CBS) in kilobytes (kbytes). On 3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928, 3942, 5142 and 5160 switches, the CBS entered value, ranging from 0-262144 Kbytes, is adjusted by rounding up depending on 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. Global parent profiles are single meters and can be provisioned with CIR and CBS only.
ebs	Excess burst size (EBS) in kbytes. On 3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928, 3942, 5142, and 5160 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 on a reboot. This does not apply to HIM profiles.
precedence	Precedence applies to the global l2-l3 profile. This number must be unique across all global l2-l3 profiles. It determines the precedence of the profile with the lowest number being the highest precedence. When the precedence value is specified, you can configure overlapping profiles. When a packet matches multiple profiles, the precedence value determines which profile to use. Global parent profiles do not have a precedence value.

For the related procedures, refer to:

- “Enabling and disabling global profiling” on page 8-140
- “Enabling and disabling global parent profiling” on page 8-141
- “Setting the traffic profiling provisioning mode, profile mode, and profile type” on page 8-142
- “Displaying traffic profiling information in port profile mode” on page 8-144
- “Setting traffic profiling port attributes for standard profiles” on page 8-147
- “Configuring a traffic profiling standard profile” on page 8-149
- “Configuring a traffic profiling global traffic profile” on page 8-155
- “Configuring traffic profiling global profile parent” on page 8-161
- “Configuring per-port standard traffic profiling” on page 8-163
- “Configuring per-port and per-VLAN standard traffic profiling” on page 8-166

Copyright© 2022 Ciena® Corporation

Traffic profiling with VLAN range support

A range of VLANs can be set for the VLAN attribute in the following modes:

- the traffic profiling profile mode is set to global-l2-l3

  Or
- the traffic profiling profile mode is set to port and the traffic profiling port mode is set to advanced
- The VLAN range option is not supported for hierarchical traffic profiles in any mode.
- The VLAN range option is not supported for port mode with resource type as extended.

  Note: The VLAN attribute with up to four VLANs is supported in global-l2-l3 mode. The VLAN range can be configured with any number of continuous VLANs. If the VLAN attribute has four or a smaller number of VLANs, then traffic profile will behave as described in"Global l2-l3 profiling mode".

The option to set a range of VLANs is intended to support traffic profiles for UNI members having EVPLs with continuous CVIDs. The device uses the resolved cos instead of qualifying on the VLAN range in the hardware.

The following limitations and recommendations apply to the VLAN range option:

- It is recommended to add a UNI port as a classifier in traffic profiles having the VLAN range attribute.
- Only a single port can be added as classifier in a global profile having the VLAN range attribute.
- A standard profile having the VLAN range as a classifier cannot have dot1dpri, dscp, ip-prec and untagged as a classifier.
- A port having a standard profile with a VLAN range classifier cannot have any hierarchical standard profiles.
- It is recommended to add a VLAN classifier with the exact same range as that configured for an EVPL member.
- There should be one traffic profile corresponding to each EVPL with a VLAN range. In case two EVPLs are configured with range and both have same encap-fixed-dot1dpri value, then it is recommended to create a single traffic-profile for any of the EVPLs. If the user creates two profiles, then the metering results can be non-deterministic.
- If a traffic-profile is created with the VLAN classifier having a VLAN range, the following traffic will get metered by this profile:

Copyright© 2022 Ciena® Corporation

— All the traffic ingressing for EVPL having exact same range.

— All the priority tagged traffic if the resolved cos (based on User configuration) for these frames is also same as encap-fixed-dot1dprifixed-dot1d-pri value set for EVPL with same CVID range.

— All the EVPLs (with single CVID or range of CVIDs) for which the resolved cos (based on User configuration) is same as encap-fixed-dot1dpri fixed-dot1d-pri value set for EVPL with same CVID range and there is no specific traffic profile configured for EVPLs with single CVID.

VLAN range support only applies when traffic-profiling is used for UNI ports of both Q-in-Q and MPLS-based Virtual Switches.

For the related procedures, see:

- "Setting the VLAN range for standard traffic profiles"
- "Configuring standard traffic profiling for EVPLs with the vlan-range-type as group"
- "Configuring a global profile with the VLAN range option"

Traffic profiling with hierarchical ingress metering

On the 3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928, 3942, 5142, and 5160 platforms, traffic profiling supports hierarchical ingress metering (HIM). When HIM is configured, the child profiles configured with a common parent profile are metered hierarchically with the parent meter. The main purpose of the parent meter is to limit the overall shared bandwidth that can be used by the child meters while allowing each child to use up to the maximum parent bandwidth. HIM is supported in global-l2-l3 profile mode. In global-l2-l3, simply create a parent profile and then create child profiles associated with the parent. HIM is supported, in profile mode, by setting the per-port traffic profiling mode to one of these 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.
- Advanced - supports both hierarchical and non-hierarchical traffic. See Advanced classification mode.

Copyright© 2022 Ciena® Corporation

Note: 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 the EVC advanced mode classification figure. Advanced mode is the only mode that supports this type of EVC based classification.

When configuring EVC-based traffic-profiling, these 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.

Figure 8-3
EVC advanced mode classification

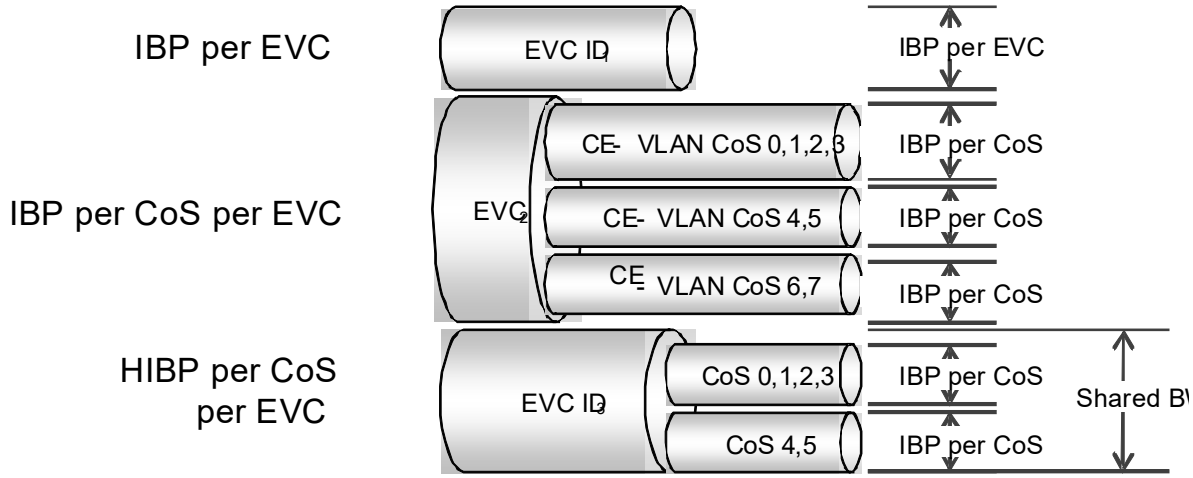

Copyright© 2022 Ciena® Corporation

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 can 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 standard traffic profiles" on page 8-169
- "Configuring hierarchical standard port traffic profiles" on page 8-175
- "Configuring VS classification for standard traffic profiles" on page 8-177
- "Configuring VS classification for standard HIM traffic profiles" on page 8-187

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.

When changing to or from a hierarchical mode, all traffic profiles 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.

392X Traffic profile scale

Traffic-profiles are configured at the edge of the network to ensure customer traffic is metered before entering the provider network. Based on the number of customers, a large number of traffic-profiles can be required.

Copyright© 2022 Ciena® Corporation

The scale of traffic profiles depends upon the available hardware resources on the device. To support more scale, more hardware resources are required and can be obtained by using a combination of different hardware resource blocks.

This feature is supported in the ‘port’ traffic-profiling mode.

392X scale is 744 default type, 792 extended type, for a total of 1536.

Standard Traffic Profile (Port Mode)

Standard traffic profiles will be of two different types based on hardware resources they consume:

- default type
- extended type

The default resource-type traffic-profile supports a wide range of classifier sets and can be created with either statistics ON or OFF. This type of traffic-profile is supported in all types of supported port-mode (standard-dot1dpri, standard-vlan, advanced, etc.). In other words, all the existing traffic-profiles are converted to “default” resource-type traffic-profiles and inherit the default resource-type traffic-profile properties.

The extended resource-type traffic-profile is similar to the default resource-type traffic-profile however, supports only a limited classifier sets:

- (port + VLAN) or (port + VS + VLAN)
- port + VS

The port can be either an EPL/EVPL member of a virtual-switch or an MPLS AC port.

Extended resource-type traffic-profiles are only supported in a “standard-vlan” and in “advanced” port-mode. Hierarchal metering classifiers are not supported using extended type traffic-profile. When an extended type traffic-profile is created, the statistics are ON and can not be disabled.

Before creating extended resource-type traffic-profile, the port must be configured as UNI or as an AC port. When “port +vlan” is set as the classifier, the vlan represents the customer-vlan (CVID). In addition, if the vlan is not added as an EVPL member of any virtual-switch for the port, then the “extended” resource-type traffic profile cannot be created for that vlan.

For example, if a port is added as an EPL member and the traffic-profile is created using “port + vlan” as the classification and the configured “vlan” is not added as an EVPL member, can only be assigned as a “default” resource-

Copyright© 2022 Ciena® Corporation

type. Attempts to assign as “extended” resource-type are denied. If the extended type traffic-profile is created on any port, then that port cannot be removed from the last EVPL member or EPL member of the virtual-switch.

Extended resource-type traffic-profiles support the same actions and statistics as supported by the default traffic-profile. Supported actions are:

- remark-color-policy
- remark-rcos-policy
- dscp-remark-policy

Both types of traffic-profiles can co-exist on the same port however, the classification criteria should not overlap. Once a traffic-profile is created, the resource-type cannot be modified.

SAOS software uses the below algorithm to determine the resource-type of the traffic-profile need to be created to get the maximum scale of traffic-profiles.

Copyright© 2022 Ciena® Corporation

Figure 8-4
Flow Chart to Determine Traffic-profile Resource Type

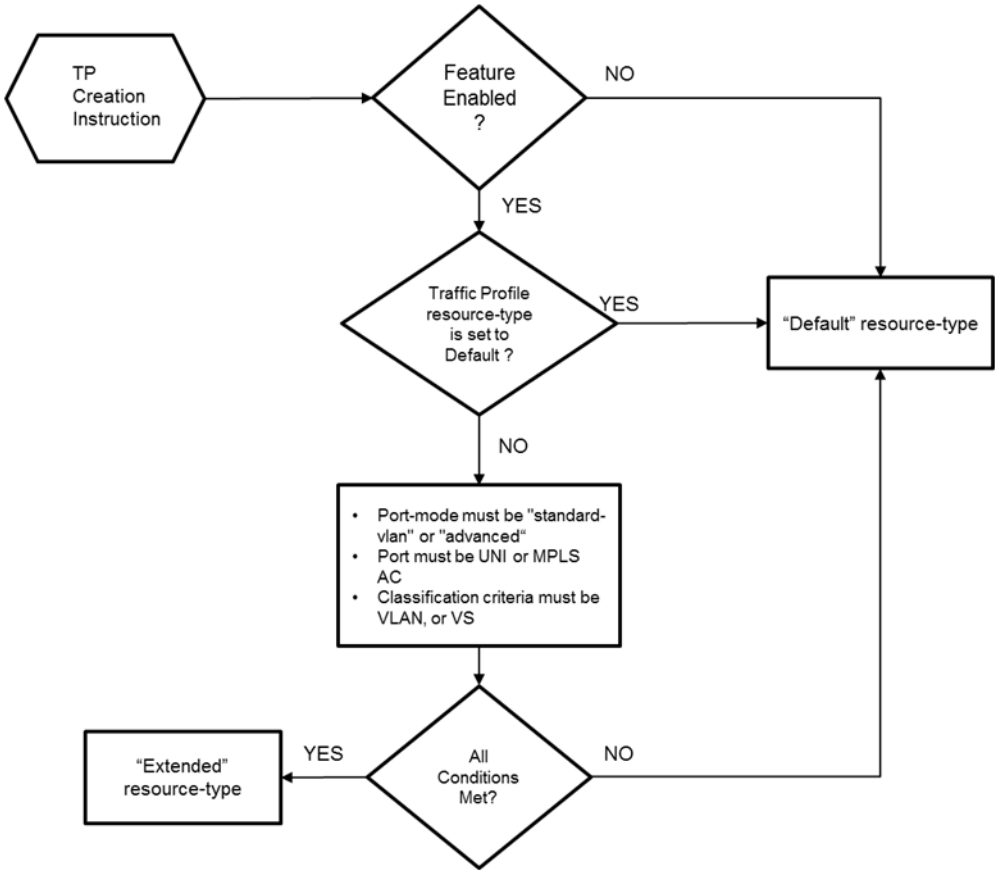

Extended type traffic-profile caveats

- For extended resource-type traffic-profiles, meter bandwidth calculations for CIR and EIR/PIR do not consider IFG (Inter Frame Gap). There is no impact configuring bw-calculation-mode to “transport”. By default, the calculation mode is configured as “payload” where IFG is not considered and extended type traffic-profiles work correctly.
- Untagged classifier and CVID classifier are not be supported when using extended resource-type traffic-profile.
- Extended resource-type traffic-profile only support “narrow” classification mode and there is no change in behavior when the classification mode is changed to “wide”.
- An extended type traffic-profile cannot be created on a port when a default type traffic-profile exists on the port for without an EVPL member added.
- If an extended resource-type traffic-profile exists on the port, the “default” type traffic-profile creation is not allowed if no EVPL member added.

Copyright© 2022 Ciena® Corporation

Resource manager

A resource manager is required for “extended” resource-type traffic-profiles. Hardware resources used are shared among other features. The resource manager ensures that resources are shared between the different features.

Default type traffic-profiles use the existing resource manager for allocation of classifiers, meters and counters under the traffic-profile feature. Allocation of resources for this type of traffic-profile follow the existing approach.

The resource manager for extended resource-type traffic-profile is dynamic in nature. This means that the resources are shared among multiple features without pre-allocating the resources. Resources are allocated on a first-come, first-served basis.

MEF 10.3 ingress metering

MEF 10.3 defines a bandwidth profile algorithm to process more than one flow and allow token sharing among the flows. A bandwidth profile is specified using bandwidth profile flow and the envelope.

The bandwidth profile flow is a set of service frames arriving at a UNI port that meets special criteria. This list outlines the specific criteria defined in MEF 10.3 for ingress bandwidth profile:

- All ingress service frames at the UNI
- All ingress service frames at the UNI that are mapped to a given EVC
- All ingress service frames at the UNI that have a given Class of Service and mapped to a given EVC

Copyright© 2022 Ciena® Corporation

This table describes bandwidth flows and envelopes.

Table 8-10
Bandwidth flows and envelopes

Item	Description
Bandwidth flow	• Unique name and rank in the envelope • Same name in different envelopes, but not necessary • Default name is ENVELOPE_NAME.RANK • Supports color-aware (default) and color-blind mode • Supports running counters to count accepted and dropped packets ***Note:*** The accepted counter counts green and yellow color packets. The dropped counter counts red color packets.
Envelope	• Set of bandwidth profile flows • Unique rank assigned to these bandwidth profile flows in the envelope to determine how tokens are shared among the flows • Unique name and is not the same name as the configured standard or global traffic profiles on the device

Ingress traffic mapped to a RCoS or PCP that does not correspond to a RCoS or PCP bandwidth profile flow in the envelope attached to that port, passes through without being metered.

The sharing of bandwidth can also be achieved using hierarchical ingress metering, but this type of bandwidth sharing is not predictable. MEF 10.3 sharing is more predictable and controlled using the coupling flag parameter.

These figures show different kinds of envelopes and bandwidth flows.

The Bandwidth profile at UNI port figure shows a single bandwidth profile flow created at the UNI port. All the frames ingressing to the UNI port are metered using the associated bandwidth profile. There is no sharing of tokens as the single flow is defined per UNI port.

Copyright© 2022 Ciena® Corporation

Figure 8-5
Bandwidth profile at UNI port

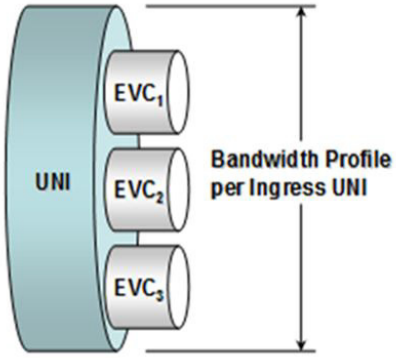

The Bandwidth profile based on EVC figure shows bandwidth profile flows based on EVC. Multiple EVCs can be added to an envelope. Bandwidth sharing occurs among these EVCs. Multiple envelopes can exist on the same UNI port, but the same EVC cannot be added to different envelopes on a single UNI port.

Figure 8-6
Bandwidth profile based on EVC

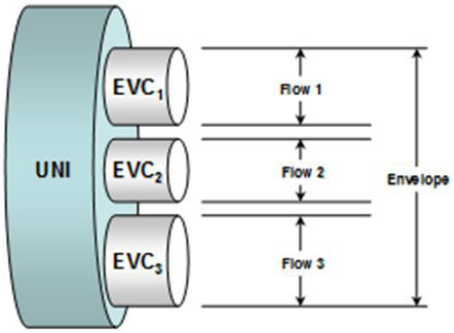

The Bandwidth profile based on CoS figure shows bandwidth profile flows based on CoS mapped to the EVC. Bandwidth sharing is done among flows.

Copyright© 2022 Ciena® Corporation

Figure 8-7
Bandwidth profile based on CoS

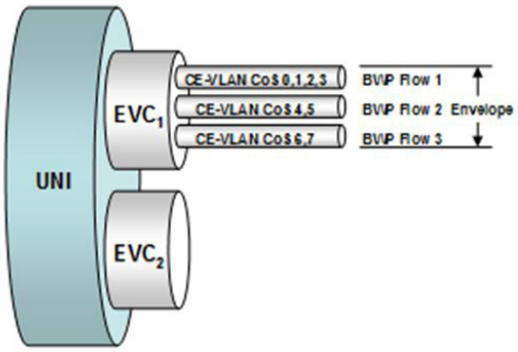

Envelope

Each envelope has these parameters:

- Envelope Id
- Coupling Flag

The envelope coupling flag allows the sharing of unused tokens from the lowest rank committed bucket to the highest rank excess bucket. The envelope coupling flag is 0 when only single bandwidth profile flow is added to the envelope.

> ***Note:*** Setting the envelope coupling flag to ON is not supported in this release. By default, the envelope coupling flag is set to OFF. This parameter cannot be changed by the user.

This table shows the number of envelopes supported by each platform.

Table 8-11
Envelopes supported by platform

Platform	Number of envelopes supported	Number of envelopes supported with statistics turned on
3924, 3926, 3928	510	62
3942	2046	762
5142, 5160	4092	4092

Envelope ingress metering is supported for Ethernet (QinQ) transport network.

Bandwidth profile parameters

Each bandwidth profile flow in the envelope has these parameters:

- CIR

Copyright© 2022 Ciena® Corporation

- CBS
- CIRmax
- EIR
- EBS
- EIRmax
- Coupling flag
- Color mode
- Rank

These parameters can also be modified when the envelope is administratively enabled, but traffic is impacted.

CIRmax

The characteristics of CIRmax are as follows:

- CIRmax limits the rate of tokens added to the committed bucket from higher rank flows.
- Bandwidth profile flow does not use overflow tokens of the committed bucket from higher rank flow in an envelope when CIRmax equals CIR. The same applies to excess bandwidth.

  Note: The default value is equal to CIR. CIRmax does not exceed the maximum device speed, and it does not depend on the port admin speed.

CIRmax is a function of CIR. The configured value is adjusted to its normalized value based on the normalized CIR value. The CIRmax value in the CIR max and CIR ratio table is expanded until 8 * CIR. Beyond this value, CIRmax does not limit the overflow tokens from higher rank flows or no limit is applied. In this case, the flow CIR is limited by the aggregate CIR bandwidth of the higher rank flows.

Table 8-12
CIR max and CIR ratio

CIRmax/CIR ratio	CIRmax value
1.00 <= r <= 1.12	1.00 * CIR
1.13 <= r <= 1.37	1.25 * CIR
1.38 <= r <= 1.62	1.50 * CIR
1.63 <= r <= 1.87	1.75 * CIR
1.88 <= r <=2.12	2.00 * CIR
2.13 <= r <=2.37	2.25 * CIR

Copyright© 2022 Ciena® Corporation

Table 8-12
CIR max and CIR ratio

CIRmax/CIR ratio	CIRmax value
2.38 <= r <= 2.62	2.50 * CIR
2.63 <= r <= 2.87	2.75 * CIR
2.88 <= r <= 3.12	3.00 * CIR

EIRmax

The EIRmax limits the tokens added to the excess bucket. Bandwidth profile flow does not use overflow tokens of the committed bucket from a higher rank flow in an envelope when CIRmax is equal to CIR. The same is true for excess bandwidth. The EIR max value of the bandwidth profile flow is greater than or equal to EIR.

> ***Note:*** The default value is equal to EIR. EIRmax cannot exceed the maximum device speed and does not depend on the port admin speed.

EIRmax is a function of EIR and follows Table 8-12 and behavior as CIRmax.

CIR and EIR relationship

The CIR/EIR value of the bandwidth profile flow is greater than or equal to 0, but for any flow, CIR/EIR/CIRmax/EIRmax is not 0.

> ***Note:*** The default value of CIR/EIR is 0. CIR/EIR cannot exceed the maximum device speed and does not depend on the port admin speed.

All the values are normalized using step size. Step size for CIR/EIR is shown in this table.

Table 8-13
CIR and EIR

CIR/EIR minimum value (Kbps)	CIR/EIR Max Value (Kbps)	Step size (Kbps)
1	4194304	8
4194305	8388608	16
8388609	16777216	32
16777217	33554432	64
33554433	67108736	128

Copyright© 2022 Ciena® Corporation

The step size for rate depends on the burst size value. Small rate and large burst size has more step size for rate. Excess and committed follow the same step size. If the excess rate is greater than the committed rate, then the step size is based on the excess rate.

CBS and EBS relationship

When CIR is greater than zero, CBS is greater than zero. When EIR is greater than 0, EBS is greater than zero.

Note: The minimum configured or default value of CBS/EBS is 4 Kbytes.

All values are normalized using step size. Step size for CBS/EBS is shown in this table in Kbits whereas allowed configured value is in Kbytes.

Table 8-14
CIR and EBS

CBS/EBS minimum value (Kbits	CBS/EBS Max Value (Kbits)	Step size (Kbits
32	16777	4
16778	33554	8
33555	67108	16
67109	134217	32

Burst step size also depends on rate values.

Bandwidth profile flow coupling flag

The bandwidth profile flow coupling flag parameter allows the sharing of overflow tokens from the committed bucket to the excess bucket in color-aware mode. Rank parameter defines the ranking of the bandwidth profile flow among other bandwidth profile flows in the same envelope. If there are “n” bandwidth profile flows in the envelope, then each bandwidth profile flow is assigned a unique rank between 0 (lowest) and n-1 (highest). The default is 0.

Classifiers and actions

The list of supported classifiers for envelope bandwidth profile flow are as follows:

- VS
- VS + RCoS (resolved CoS)
- VS + CE-VLAN priority (inner tag pcp)

An envelope can be created using any of these classifiers and then attached to the UNI port. Each bandwidth profile flow must belong to a single envelope, and an ingress frame must be mapped to a single bandwidth profile flow. An

Copyright© 2022 Ciena® Corporation

envelope must have all the bandwidth profile flows with the same classification criteria. An envelope profile flow classification mode is always narrow. The envelope bandwidth profile flow classification mode does not depend on the traffic profile classification mode. If the classifier is not defined for an envelope flow, then the bandwidth profile is done on all the traffic ingressing from the port associated with the envelope.

A UNI port cannot be removed from a virtual switch completely if the envelope attached to that UNI port and envelope is enabled. At least one member of that UNI port must be present for that virtual switch, either EVPL member or EPL.

Untagged packets are metered by CE-VLAN priority 0 meter, in case of VS+CE-VLAN priority classifier.

You cannot delete a virtual switch if the bandwidth profile flow is created for the virtual switch even if the envelope is disabled.

Traffic above EIR or frames marked RED by the meter are dropped.

This table describes some examples of not allowed classification configurations.

Table 8-15
Classification configurations

Scenario	Classifier	Action
1	Envelope-1: VS1, VS2, VS3 + RCoS-2 -> Port-X	In this envelope, 2 flows are based on the VS, but the third flow is based on the VS and RCos. This envelope does not fulfill the requirement of the same classification criteria in an envelope.
2	Envelope-1: VS1 + RCoS-1, VS2+ RCoS-3 -> Port-X	In this envelope, both the flows don't have the same VS. This configuration is not allowed because of an SAOS 6.x limitation.
3	Envelope-1: RCoS-0, RCoS-1 -> Port-X	No VS is defined. This classification criteria are not allowed as per MEF 10.3 requirements.
4	Envelope-1: VS1, VS2 -> Port-X	Both envelopes share the same VS1 on port X, so the envelopes have overlapping bandwidth profile flows that are not allowed.
	Envelope-2: VS1, VS3 -> Port-X	

Copyright© 2022 Ciena® Corporation

Table 8-15
Classification configurations

Scenario	Classifier	Action
5	Envelope-1: VS1, VS2 -> Port -X	Envelope-1 and Envelope-2 share VS1 on Port-X so envelopes have overlapping bandwidth profile flows that are not allowed.
	Envelope-2: VS1 + RCoS1, VS1 + RCoS2 -> Port X	
6	Envelope 1: VS1 + RCoS1, VS1 + RCoS2 -> Port X	This configuration is not allowed because of an SAOS 6.x limitation.
	Envelope-2: VS1 + RCoS3, VS1 + RCoS4 -> Port -X	

Frames that are marked red by the meter are dropped. The color of the frames is remarked by the meter for green and yellow color frames.

Use case 1

The Use case 1 figure shows three bandwidth profile flows configured with CIR and CIRmax. As EIR and EBS are configured as 0, this use case is suitable for color-blind mode.

Note: Since flow CoS-H is the highest rank member in the envelope, its specified CIR max does not affect the function of the envelope.

Copyright© 2022 Ciena® Corporation

Figure 8-8
Use case 1

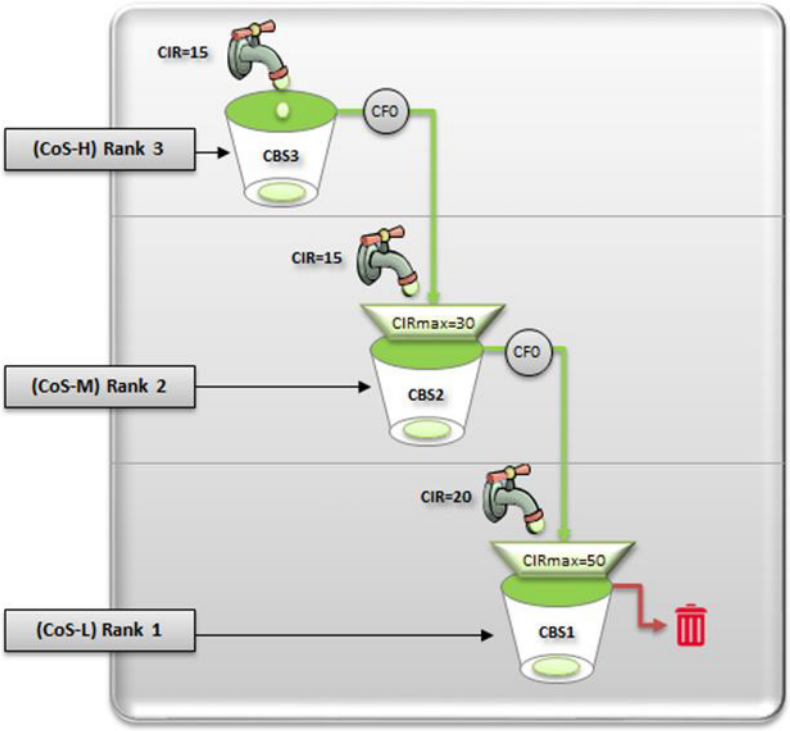

This table shows the input rates and meter output rates for different scenarios for this use case's configuration.

Table 8-16
Use case 1: Input rates and meter output rates

Scenario Number	Input rate			Meter output rate		
	H	M	L	H	M	L
1	50	50	50	15	15	20
2	0	50	50	0	30	20
3	0	0	50	0	0	50
4	50	0	50	15	0	35
5	0	15	50	0	15	35
6	5	50	50	5	25	20
7	5	5	50	5	5	40

Copyright© 2022 Ciena® Corporation

Use case 2

The Use case 2 figure shows three bandwidth profile flows that are configured with CIR and CIRmax. EIR and EBS are configured as 0. CIRmax equals CIR for flow-M so flow-M does not use bandwidth of flow-H. Flow-L can use unused bandwidth of flow-H and flow-M. The coupling flag of each flow and coupling flag of envelope are configured as 0. As EIR and EBS are configured as 0, this use case is more suitable for color-blind mode.

Note: Since flow CoS-H is the highest rank member in the envelope, its specified CIR max does not affect the function of the envelope.

Copyright© 2022 Ciena® Corporation

Figure 8-9
Use case 2

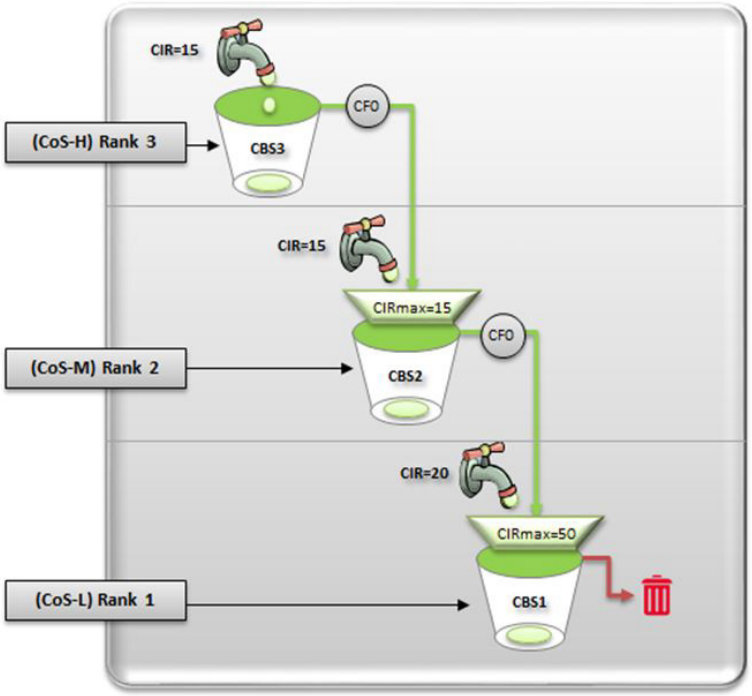

This table shows the input rates and meter output rates for different scenarios for this use case's configuration.

Table 8-17
Use case 2: Input rates and meter output rates

Scenario Number	Input rate			Meter output rate		
	H	M	L	H	M	L
1	50	50	50	15	15	20
2	0	50	50	0	15	35
3	0	0	50	0	0	50
4	50	0	50	15	0	35
5	0	15	50	0	15	35
6	5	50	50	5	15	30
7	5	5	50	5	5	40

Copyright© 2022 Ciena® Corporation

Use case 3

The Use case 3 figure shows three bandwidth profile flows configured with CIR and CIRmax. EIR and EBS are configured as 0. The coupling flag of each flow and coupling flag of envelope are configured as 0. As EIR and EBS are configured as 0, this use case is suited for color-blind mode.

Note: Since flow CoS-H is the highest rank member in the envelope, its specified CIR max does not affect the function of the envelope.

Figure 8-10
Use case 3

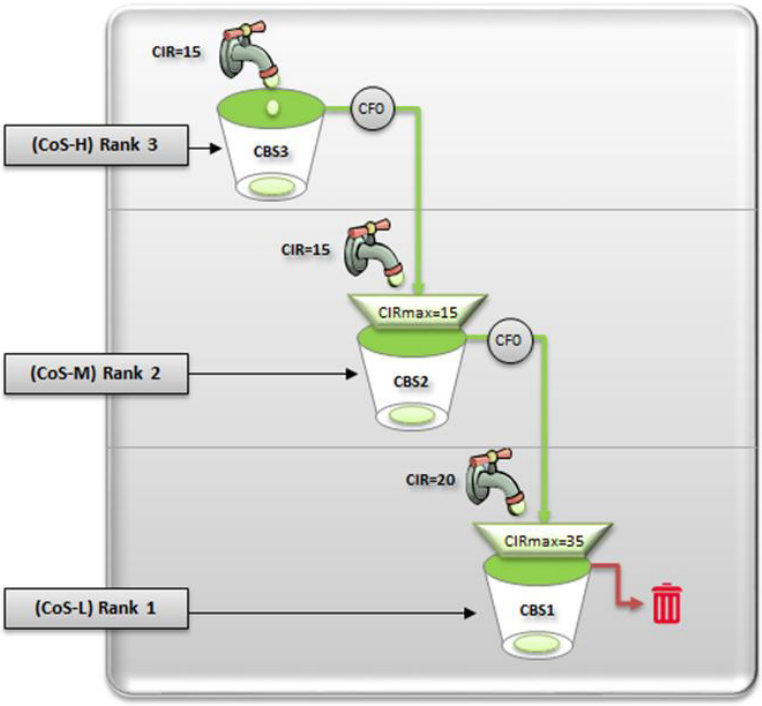

This table shows the input rates and meter output rates for different scenarios for this use case's configuration.

Table 8-18
Use case 3: Input rates and meter output rates

Scenario Number	Input rate			Meter output rate		
	H	M	L	H	M	L
1	50	50	50	15	15	20
2	0	50	50	0	15	35
3	0	0	50	0	0	35

Copyright© 2022 Ciena® Corporation

Table 8-18
Use case 3: Input rates and meter output rates

Scenario Number	Input rate			Meter output rate		
	H	M	L	H	M	L
4	50	0	50	15	0	35
5	0	15	50	0	15	35
6	5	50	50	5	15	30
7	5	5	50	5	5	35

Use case 4

The Use case 4 figure shows three bandwidth profile flows configured with committed and excess bucket parameters (CIR, CIRmax, EIR, EIRmax, CBS and EBS). The coupling flag of each flow and the coupling flag of the envelope are configured as 0. As excess bucket is also configured, this use case can work in color-aware and color-blind mode.

Unused tokens from the committed bucket are used by the committed bucket of lower rank bandwidth profile flow and unused tokens from the excess bucket are used by the excess bucket of lower rank flow. CIRmax and EIR max limit the rate of tokens added to the committed and excess bucket.

> ***Note:*** Since flow CoS-H is the highest rank member in the envelope, its specified CIR max does not affect the function of the envelope.

Copyright© 2022 Ciena® Corporation

Figure 8-11
Use case 4

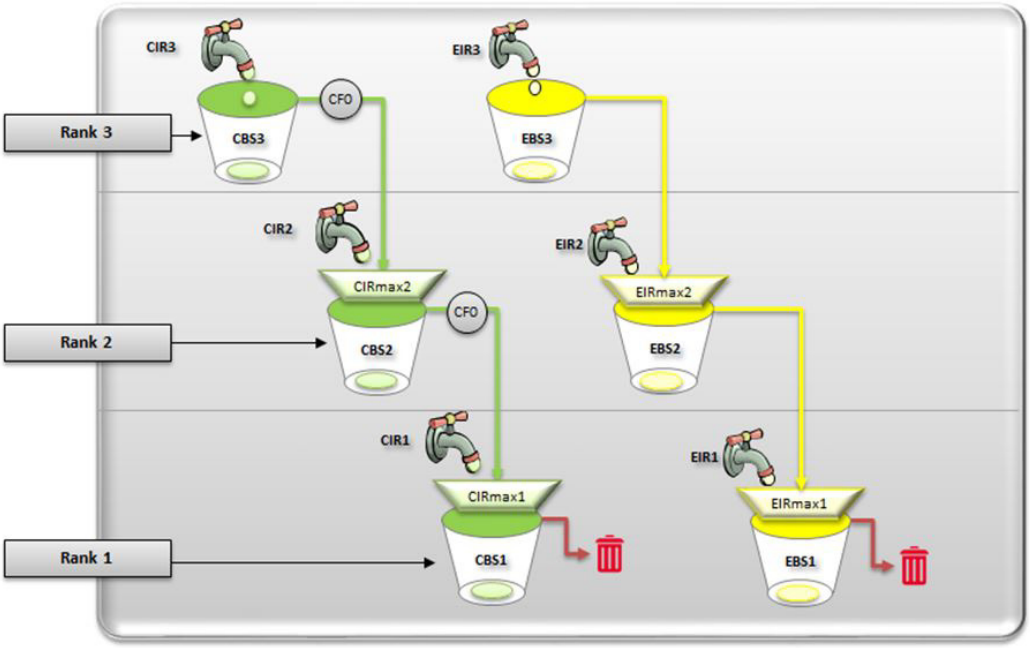

The configuration values for this use case are as follows:

- Cos-H:
  - CIR=15Mbps
  - EIR=10Mbps
- COS-M:
  - CIR=15Mbps
  - CIRmax=15Mbps
  - EIR=10Mbps
  - EIRmax=20Mbps
- COS-L:
  - CIR=20Mbps
  - CIRmax=35Mbps
  - EIR=10Mbps
  - EIRmax=30Mbps

Copyright© 2022 Ciena® Corporation

This table shows the input rates and meter output rates for different scenarios for this use case's configuration.

Table 8-19
Use case 4: Input rates and meter output rates

Scenario Number	Input rate			Meter output rate (green)			Meter output rate (yellow)		
	H	M	L	H	M	L	H	M	L
1	50	50	50	15	15	20	10	10	10
2	0	50	50	0	15	35	0	29	10
3	0	0	50	0	0	35	0	0	15
4	50	0	50	15	0	35	10	0	15
5	0	15	50	0	15	35	0	0	15
6	5	50	50	5	15	30	0	20	10
7	5	5	50	5	5	35	0	0	15
8	80	80	80	15	15	20	10	10	10
9	15	80	80	15	15	20	0	20	10
10	15	15	80	15	15	20	0	0	30
11	80	15	80	15	15	20	10	0	20
12	15	25	80	15	15	20	0	10	20
13	20	80	80	15	15	20	5	15	10
14	20	20	80	15	15	20	5	5	20

Use case 5

The Use case 5 figure shows three bandwidth profile flows configured with committed and excess bucket parameters (CIR, CIRmax, EIR, EIRmax, CBS and EBS). The coupling flag of each flow is 1 and the coupling flag of the envelope is 0. EIR is configured as EIRmax, so that overflow tokens of the excess bucket of a higher rank flow is not used by the excess bucket of the lower rank flow. The CIRmax parameter does not limit the overflow tokens from the committed token bucket to the excess token bucket when the coupling flag is enabled for the bandwidth profile flow.

Overflow tokens from the committed bucket are used by the excess bucket of the same rank bandwidth profile flow and unused tokens from the excess bucket are dropped.

Copyright© 2022 Ciena® Corporation

Note: Since flow CoS-H is the highest rank member in the envelope, its specified CIR max does not affect the function of the envelope.

Figure 8-12
Use case 5

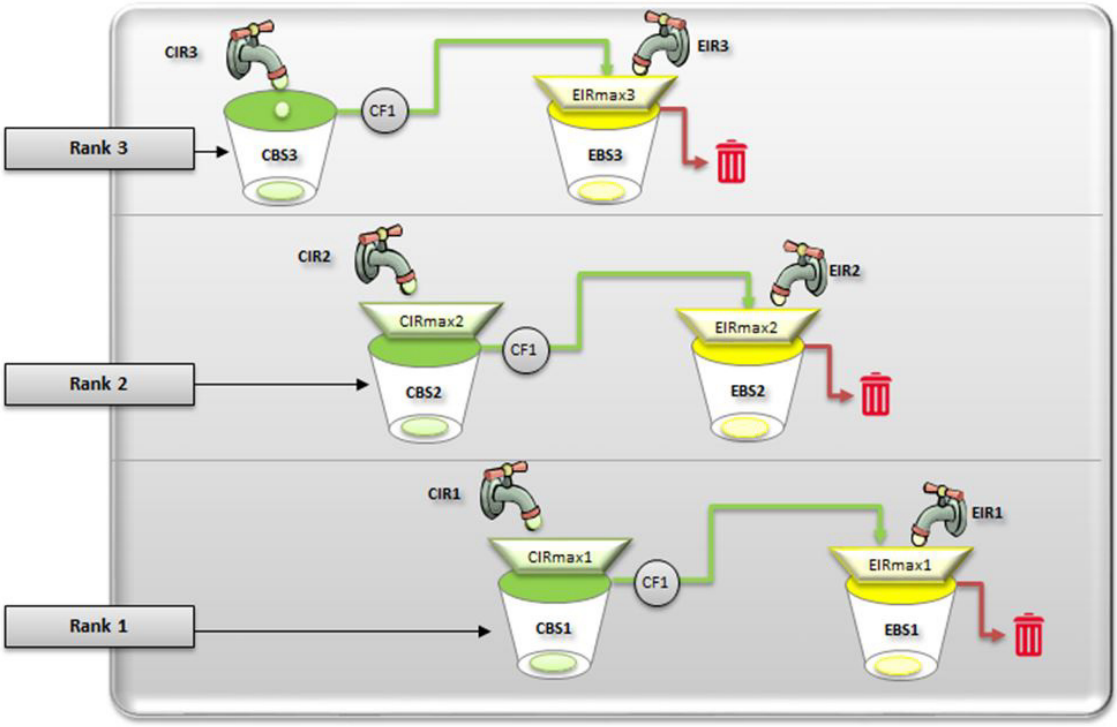

The configuration values for this use case are as follows:

- Cos-H:
  - CIR=15Mbps
  - EIR=10Mbps
- COS-M:
  - CIR=15Mbps
  - CIRmax=15Mbps
  - EIR=10Mbps
  - EIRmax=10Mbps
- COS-L:
  - CIR=20Mbps
  - CIRmax=20Mbps
  - EIR=10Mbps

Copyright© 2022 Ciena® Corporation

— EIRmax=10Mbps

This table shows the input rates and meter output rates for different scenarios for this use case's configuration.

Table 8-20
Use case 5: Input rates and meter output rates

Scenario Number	Input rate (green)			Input rate (yellow)			Meter output rate (green)			Meter output rate (yellow)		
	H	**M**	**L**	**H**	**M**	**L**	**H**	**M**	**L**	**H**	**M**	**L**
1	50	50	50	50	50	50	0	0	0	15	15	20
2	0	50	50	0	0	0	0	15	20	0	10	10
3	0	0	50	0	0	0	0	0	20	0	0	10
4	50	0	50	0	0	0	15	0	20	10	0	10
5	0	15	50	0	0	0	0	15	20	0	0	10
6	5	50	50	0	0	0	5	15	20	0	10	10
7	5	5	50	0	0	0	5	5	20	0	0	10
8	50	50	50	40	40	40	15	15	20	10	10	10
9	15	50	50	0	40	40	15	15	20	0	10	10
10	15	15	50	0	0	40	15	15	20	0	0	10
11	50	15	50	40	0	40	15	15	20	0	0	10
12	15	15	50	0	10	40	15	15	20	0	10	10
13	15	50	50	5	40	40	15	15	20	5	10	10
14	15	15	50	5	5	40	15	15	20	5	5	10

Envelope statistics

Envelope statistics are supported per bandwidth profile flow basis. Enabling and disabling envelope statistics is provided at the envelope level. Statistics are enabled or disabled for all bandwidth profile flows in the envelope, but they cannot be controlled at the per bandwidth flow level. Statistics count accepted and dropped frames. An accepted packet statistics count green and yellow packets marked by the meter.

Bandwidth profile flow statistics are displayed in packets and in bytes. By default, statistics is set to ON. If all the hardware resources for counters are consumed, then the envelope can be created with "statistics off" or you can set the statistics to OFF for already created envelopes to continue to create new envelopes with statistics on.

Copyright© 2022 Ciena® Corporation

Statistics for a bandwidth profile flow can be obtained using SNMP for:

- Accepted packets
- Dropped packets
- Accepted bytes
- Dropped bytes

If the classifier (r-cos or pcp) value is modified for bandwidth profile flow, statistics are reset to 0 before starting to collect statistics for a new classifier for the same flow. Statistics are cleared using the management interface and it is set to 0. On disabling the envelope, statistics are stopped and cleared.

Interaction with standard and global traffic profiles

This section describes the interaction with standard and global traffic profiles.

- A standard traffic profile and envelope meter cannot exist on the same UNI port. An envelope meter cannot be created if the standard profile is created on that port. A standard profile cannot be created when the envelope is enabled on that port.
- The envelope meter classifier mode is always “narrow” and has no impact when modified by “traffic-profiling set port <port-no> classifier-mode” CLI, in port mode.
- The envelope ingress meter functionality is independent of “arp-standard-profile” actions when the traffic profile mode is set to “port”.
- When envelope metering is used, and standard-profiles are not being used on a particular port, then the mode on that port must be set to “none” or “advanced”. If the port mode is set to other than the suggested modes, then traffic may be dropped because of this non-conform entry.
- The traffic-profiling admin state on the port has no impact on the envelope meter’s operational state, whereas if the port mode is not set to “none” or “advanced”, then traffic may be dropped where the traffic-profile admin state on that port is set to enabled. It is recommended that the port mode is set to “none” or “advanced” when using envelope metering.
- The global traffic profile cannot be created for all ports if the envelope is already created in global-l2-l2 mode and vice-versa, envelope cannot be created if the global traffic profile is created on all ports even if the global traffic profile or envelope is in disable state.
- Ports cannot be added to the global traffic profile if the envelope is created on that port. The envelope can also not be attached to a port if global traffic profile is created on that port even if global traffic profile or envelope is in disable state.

Copyright© 2022 Ciena® Corporation

- Bandwidth profile flow uses the EIR and EIRmax parameters when the traffic profile configuration mode is set to “EIR” and convertw all EIR and EIRmax values to PIR and PIRmax values when changing mode to “PIR”. PIRmax is the sum of CIRmax and EIRmax. When setting PIR less than CIR, then PIR is set equal to CIR.

Limitations

This section describes the limitations for MEF 10.3 ingress metering.

- A maximum of eight bandwidth profile flows can be added to the envelope.
- The list of CE_VLAN priority and RCoS are not supported as classifier.
- On changing the r-cos or dot1dpri (c-pcp), the rank automatically changes because the rank is equal to the r-cos or dot1dpri value.
- The coupling fag value of all bandwidth profile flows in an envelope must be the same.
- EIRmax must be equal to EIR when the coupling flag of a bandwidth profile flow is set to on.
- Traffic is impacted when changing the envelope or bandwidth profile flow parameters when the envelope is enabled.
- MPLS VS and sub-port VS as a classifier is not supported.
- Envelopes can be attached to UNI ports only.
- L2 control frames are subject to meter by envelope bandwidth profile flow. Frames marked as drop by the envelope meter can be further redirected to CPU port or some other egress port.
- SOAM frames (such as CCM, LM and DM) are not subject to meter by envelope meters.
- Drop action of the envelope bandwidth profile flow has a higher precedence over the allow action of ACL.
- If any traffic is classified as envelope meter, but further dropped because of ACL or Copy to CPU action, then accepted counters of the envelope meter are not incremented.
- Broadcast containment and envelope meters on the same port do not work as expected. The output rate can be the result of double metering.
- Modifying the envelope and envelope flow name is not supported.
- Modifying the VS as a classifier for envelope flow is not supported.
- When the ratio of CIRmax and CIR is greater than 8.00, envelope flow CF=0 and CIRmax is less than the sum of CIR of higher rank flows, then the CIRmax parameter does not limit the shared traffic. It shares all unused committed traffic from higher rank flows. The same limitation applies to excess traffic.

Copyright© 2022 Ciena® Corporation

- If the same port, VS and different CoS flows are mapped to different envelopes, then those envelopes can have only a single flow. If an envelope has more than one flow, then another envelope having flow with the same port, VS and different CoS cannot be enabled.
- When envelope flow classification is based on RCoS, then all RCoS related limitations for a QiQ network are applied for RCos derivation for envelope meters.
- On the 3924, 3926 and 3928 devices, envelope meters do not consider IFG for all metering calculations, irrespective of bandwidth calculation mode setting on the device.
- Envelope meters do not work with VC transform (on the NNI port) for “known unicast” traffic. However, envelope meters do work for some time after a reboot until traffic has been learned. The envelope meter statistics for accepted packets/bytes are incremented, but dropped packet/byte counters are not incremented.
- The CIR/CBS buffer is not overflowing correctly within an envelope. CIRmax is not limiting the maximum CIR of the bandwidth profile flow in the envelope in some specific configurations. The issue is seen when the CIRmax of the second highest flow is set for less than the aggregate CIR of the higher ranked flows and traffic for the highest ranked flow is less than the configured CIR.

  For example, if an envelope is configured with three flows (H, M, and L) with the following configuration and ingress traffic:

  Flow H - EVC 1000 - CIR 20M, CIRmax 20M - Ingress rate 10M
  Flow M - EVC 1001 - CIR 30M, CIRmax 37.5M - Ingress rate 100M
  Flow L - EVC 1002 - CIR 20M, CIRmax 70M - Ingress rate 100M

  The actual egress rates for H, M, and L flows are 10M, ~34M, and ~26M, whereas the expected egress rates are 10M, 37.5M, and 22.5M, respectively.

See these procedures:

- “Configuring an envelope” on page 8-120
- “Configuring a bandwidth profile flow in an envelope” on page 8-121
- “Clearing statistics for envelope flows” on page 8-123
- “Deleting an envelope” on page 8-124
- “Enabling an envelope flow” on page 8-125
- “Disabling an envelope flow” on page 8-126
- “Removing an envelope flow” on page 8-127
- “Modifying bandwidth flow attributes” on page 8-128
- “Displaying envelopes” on page 8-129

Copyright© 2022 Ciena® Corporation

- "Setting default values for envelope and flow attributes" on page 8-132
- "Configuration example for MEF 10.3 ingress metering" on page 8-133

Congestion management

Congestion management processing determines whether a frame is 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 these 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 on 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 sWRED profiles on the 3903, 3903x, 3904, 3905, 3924, 3926, 3928, 3942, 5142, and 5160

The 3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928, 3942, 5142, and 5160 platforms support sWRED profiles to handle two types of traffic:

- Green traffic
- 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

Copyright© 2022 Ciena® Corporation

- 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 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. This table describes how the default sWRED profile (Default-S-WRED) is configured.

Table 8-21
sWRED profile configuration

Thresholds and probabilities	Value
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 is the same as the Default-S-WRED profile.

Copyright© 2022 Ciena® Corporation

Egress queuing and scheduling

This section describes a port's egress queues. The description includes:

- Queue structure
- Queue scheduling
- Queue mapping
- Shaping

The egress queues are structured in two ways depending on the device:

- Non-hierarchical
- Hierarchical

Non-hierarchical egress shaping

In this queue structure, each port is allocated eight egress queues. These eight egress queues are attached to a port scheduler forming the egress port queue group. This type of egress queuing is supported on the 3903, 3903x, 3904, 3905, and 3906. By default, the queues are scheduled using the strict priority algorithm.

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 are serviced before the lower CoS queues.

In addition to strict priority, three other scheduling algorithms can be configured along with a scheduler weight. The scheduler weights determine the ratio of packets or bytes that are services from the queues. If the scheduler weight of the queue is set to 0, the scheduler treats that queue as if it is strict priority mode.

Copyright© 2022 Ciena® Corporation

This table describes the scheduling algorithms.

Table 8-22
Scheduling algorithms

Scheduler	Description
Weighted Round Robin (WRR)	Queues in the queue group are serviced to the maximum of the per queue scheduler weight based on packets per round.
Round-Robin (RR)	Queues in the queue group are services as if all the queues in the queue group have equal weights.
Weighted Deficit Round Robin (WDRR)	Queues in the queue group are serviced in a weighted fashion according to the per-queue scheduler weight and frame size. When using this 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 the port speeds lower than 10 Mbps. Small pool sizes have less latency and more jitter. Larger sizes have more latency and less jitter. The value entered is rounded to the nearest size of 50, 100, 200, 400, 800 or 1600 Kbytes. The default value is 1600 Kbytes.

The default scheduler weights for each queue are as follows:

- Queue 0: 20
- Queue 1: 30
- Queue 2: 40
- Queue 3: 50
- Queue 4: 60
- Queue 5: 70
- Queue 6: 80
- Queue 7: 0

The valid range for the scheduler weights is 0-1000.

Copyright© 2022 Ciena® Corporation

Egress queue map

At ingress, each packet is assigned an internal CoS value (RCoS) either derived from a mapping of the DSCP or the 802.1p priority value or from a fixed value. The internal CoS value is used at the egress port to determine which egress queue the packet uses. A queue map is associated with each port. This queue map is used to map the RCoS value of the packet to a queue number. The default map, Default-R-CoS, must be included in the three maps.

Egress shaping

A shaper-rate and burst-size can be configured for the egress port queue group that limits the amount of egress bandwidth (Kbps) a port uses (0 to 40 Gbps) and controls the amount of traffic (Kbytes) that can burst (0 to 262144 Kbytes). The configured shaper-rate normalizes to a rate in varying increments depending on the sum of the configured burst-size or the sum of the configured rate whichever requires the larger rounding size. This table shows the rounding steps.

Table 8-23
Rounding steps

Burst size/shaper rate	Burst/rate steps
0 Kbytes <= burst-size < 16384 Kbytes	4K/64 Kbps burst/rate steps
16384 Kbytes <= burst-size < 32768 Kbytes	8K/128 Kbps burst/rate steps
32768 Kbytes < =burst-size < 65536 Kbytes	16K/256 Kbps burst/rate steps
65536 Kbytes <= burst-size < 131072 Kbytes	32K/512Kbps burst/rate steps
131072 Kbytes <= burst-size < 524288 Kbytes	64K/102.4 Kbps burst/rate steps
0 <= shaper-rate < 16.78 Gbps	4K/64 Kbps burst/rate steps
16.78 <= shaper-rate < 33.55 Gbps	8K/128 Kbps burst/rate steps
33.55 Gbps <= shaper=rate < 80 Gpbs	16K/256 Kbps burst/rate steps

The rounding steps also pertain to the rates set for the individual queues where the burst-size equals CBS+EBS and the shaper-rate equals CIR+EIR.

to guarantee CIR, the individually-set CIR on queues must not be oversubscribed for the egress port. That is, the sum of all CIRs for the queues on a port must not exceed the administrative 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.

Copyright© 2022 Ciena® Corporation

When a queue or an attribute of the egress port 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). The shaper rate of 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 Gbyte or 10 Byte ports.

> ***Note:*** The burst size is not divided among ports in an aggregation group. The default burst size is 10240 Kbytes.

Egress queue depth

Egressing packets are enqueued onto a queue using available egress buffer space. Each port has a minimum guaranteed buffer space assigned to it. In addition, there is also a shared egress buffer space that is shared between all ports and is available for use to the ports when their minimum guarantee is fully used. Each port is allowed a configured percentage of the remaining available in the shared buffer space. By default, this percentage is set to 20% (medium). This percentage is configurable as the dynamic-max-queue-depth, as shown in this table.

Table 8-24
Dynamic-max-queue-depth percentage

CLI value	Percentage
Low	11
Medium	20
Medium-high	33
High	50

Configuration objects

The egress-port-queue-group object represents the egress queueing of the port. Each port has an egress-port-queue-group object. The egress-port-queue-group object contains these attributes:

- Shaper-rate
- burst size
- Scheduler algorithm
- Traffic shaping (on/off)
- Dynamic max queue depth

Copyright© 2022 Ciena® Corporation

Hierarchical egress queuing

This queue structure supports two levels of hierarchy which enables support for per service queuing. Per service queuing allows packets to be mapped to an egress queue based on VLAN and CoS values instead of just CoS values for non-hierarchical queuing. Hierarchical queuing is supported on the 3924, 3926, 3928, 3942, 5142 and 5160.

The lowest level of queuing hierarchy is defined by eight queues and a scheduler. The queues and scheduler form a queue group. The next level of hierarchy is defined by multiple queue groups attached to a port scheduler. This figure shows how hierarchical egress queuing works.

Figure 8-13
Hierarchical egress queuing

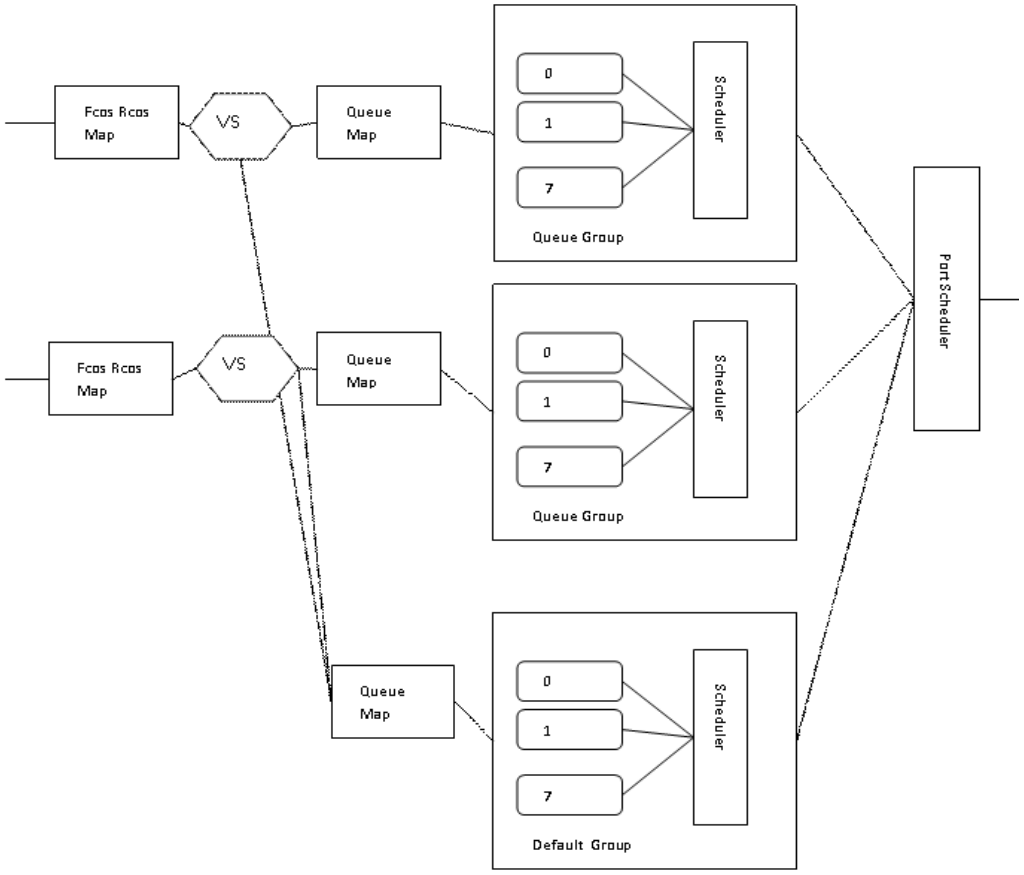

There are two types of queue groups:

- Default queue groups
- Custom queue groups

Copyright© 2022 Ciena® Corporation

Similar to the queues, the queue groups (default or custom) have configurable attributes which include scheduler weight, priority, and shaper rates.

The priority is used by the scheduler to determine the order in which to service the queue groups/queues. A higher priority gets serviced first. The priority is important when the scheduler algorithm is strict priority as the higher priority queue group/queue can block lower priority ones. With RR, WRR and WDRR, the priority is of lesser importance as each queue group/queue is serviced in the order of priority, but no queue group/queue is blocked as the queue group/queue is serviced up until its weight has been satisfied.

Due to hardware limitations on the 3924, 3926 and 3928 devices, there are restrictions on hierarchical egress queuing capability. The hardware requires that the ports pre-allocate the queue groups that are required for the device.

The limitations for the 3924, 3926 and 3928 devices include:

- The allocation of custom queue groups for each port must be specified before any custom queue groups are created
- Once a custom queue group is created, the resource allocation cannot change.
- All custom queue groups must be deleted before changing the queue group resource allocation.
- All ports attached to the custom queue group must be attached at the same priority index.
- The maximum number of custom queue groups that can be attached to a port is 7.
- There are 256 queues available for use in custom queue groups. Each queue group must contain 8 queues which means there are only 32 group instances available for the device.
- When the priority index when adding a port to a queue group, the index value can only be in the range: 0 to (queue-group-count -1).
- To enable service queuing on a VLAN/virtual switch, all member ports of a VLAN/virtual switch must have the same queue group attached when the queue group is associated to a VLAN/virtual switch in the CLI.
- An aggregation port inherits the queue group resource allocation of the port with the lowest resource allocation. When adding a port to an aggregation port, if a queue group is attached to the aggregation port, a check is made to ensure that the physical port has at least the resource allocation assigned to the aggregation port.
- When the custom queue group allocation is configured for a physical port that is already part of an aggregation port, the resource allocation for the aggregation port is recalculated.

Copyright© 2022 Ciena® Corporation

- Before adding a physical port to an aggregation port, queue group allocations must be configured on the physical port. You cannot change the queue group allocations of a port that is part of an aggregation port.

Default queue groups

Default queue groups (or egress port queue groups) are used for packets that do not match any defined service-based queues (custom queue group). There is one default queue group allocated per port. Each default queue group contains eight queues. The default queue group operates similar to the non-hierarchical egress port queue group.

By default, the default queues follow strict priority scheduling because the scheduler weights for these queues are set to zero. Even though there are multiple levels of queue hierarchy, there can only be one scheduler algorithm type used for the whole port. The scheduler is set at the port level and can be set to RR, WRR or WDRR. When the scheduler weight is set to zero, the queue or queue group is scheduled using strict priority. All queues can have their scheduler weight set to zero to indicate strict priority. Only the default queue group can be set with the scheduler weight of zero. Note that scheduler granularity cannot be set as these devices adjust the granularity based on the port speed.

The valid scheduler weight range for the default queue group is 0 (strict priority) or 10-127. The valid scheduler weight range for the default queues is 0 (strict priority) or 2-127. When determining the weights to configure, using the higher values provide for better ratio accuracy.

Multiple destination traffic always egresses out the queues of the default queue group.

Custom queue groups

A custom queue group consists of a scheduler and 8 unicast queues. Multicast queues are not supported for custom queue groups. A custom queue group is associated with a queue group profile which describes the characteristics of each of the 8 queues which includes the scheduler weight, congestion avoidance profile and the shaping rates.

A queue group profile can be associated to multiple queue groups. The rates for the queues in the profile are the configured rates which are rounded based on Table 8-23. The rounding steps may vary between ports if the EIR Policy is set to inherit as then the EIR for the queues may be different for each port and the rounding algorithm may produce a different rounding step.

The valid scheduler weight range for the custom queue group is 10-127 and the valid range for the custom queues is 0 (strict priority) or 10-127. Again, when determining the weights to configure, the use of the higher values provide for better ratio accuracy.

Copyright© 2022 Ciena® Corporation

A custom queue group is used to support service based queuing. The custom queue group can be associated with multiple ports. Each association can have different rates and scheduler weight. The association of a custom queue group with a VLAN/VS and a queue map define service based queuing for that VLAN. All the single destination traffic that contain the specified VLAN that egress out the associated ports egress out the custom queue group queues. Single destination traffic includes multicast and broadcast packets that only egresses one port. Any traffic with that VLAN that egresses out a port that is not associated with that custom queue group egress out the default queue group queues.

A maximum of 16 queue groups can be defined, but only a maximum of 6 queue groups can be associated with a VLAN or VS.

A physical port cannot be added to an aggregation if the physical port is associated with any custom queue groups. When a physical port is added to an aggregation that has associated custom queue groups, the physical port has the same custom queue group associations and characteristics as the aggregation.

This information applies to the 3924, 3926 and 3928 devices only:

- Traffic egressing out a service-enabled VLAN/VS that is not part of the service queue group egresses out the default queues.
- For VLAN/VS traffic egressing out a port that has service queuing enabled and the VLAN/VS is not service queueing enabled, then known unicast traffic egress out queues in the queue group attached at priority 0.
- Unknown unicast and multicast traffic egress out the default queues. Ports that are service queuing enabled must have a queue group attached at priority 0.
- You cannot attach more queue groups than the number that was preallocated using the traffic-service queuing resource CLI.
- The maximum number of queues that can be used for custom queue groups is 256 and each queue group must contain eight queues to give a maximum number of queue group instances of 32.
- The queue group pre-allocation must be done on the physical port prior to adding the port to an aggregate port. Preallocation cannot be done on a port that is part of an aggregate port.
- The number of pre-allocated queue groups for an aggregation port is determined by the port that has the lowest number of preallocated queue groups of the ports that are part of an aggregation.

Copyright© 2022 Ciena® Corporation

- If a port is added to an aggregate port and the port has fewer preallocated queue groups than the aggregate, then it can only be added to the aggregate port if the queue groups attached are in the lower resource allocation. Otherwise, the port cannot be added to the aggregation port.
- To enable service queuing on a VLAN/VS, a member port must be part of the queue group specified or be a port that does not have service queuing enabled. A port does not have service queuing enabled. A port does not have service queuing enabled is a port that if it is part of a queue group, that queue group is not specified for any service queuing on a VLAN/VS.
- All ports attached to the custom queue group must be attached at the same priority index.
- When setting the priority index when adding a port to a queue group, the index value can only be in this range: 0 to (queue-group-count - 1).
- The 3924, 3926 and 3928 devices must pre-allocate the number of custom queue groups that are associated with each physical port before a queue group is configured. Once the first queue group is allocated, the pre-allocation number cannot change. The pre-allocation is done through the traffic-services queuing resource CLI. the default number of queue groups pre-allocated for each port is 0.
- To change the preallocation of custom queue groups per port, all queue groups must be deleted.

Egress queue map

As described for the non-hierarchical egress queueing case the system uses the assigned internal CoS value associated to the packet to determine the egress queue using the queue map. In hierarchical egress queuing, there is a queue map specified for the port which his used for mapping to the default queues and queue maps specified for the VLAN/VS which is used for mapping the custom queues. Note that the maximum number of unique queue maps associated amongst all the VLANs/VSes is 3.

Egress shaping

Shaper rates can be configured at many levels of the hierarchy. Rates can be configured at the port, the default queue group, the port association of the custom queue group, the default queues and the custom queues.

Similar to the non-hierarchical egress queuing, the configured shaper rate normalizes to a rate in varying increments depending on the sum of the configured burst size or the sum of the configured rate whichever requires the larger rounding size. The rounding steps also pertain to the rates set for the individual queues and queue groups where the burst size is CBS+EBS and the shaper rate is CIR+EIR. The rounding steps are the same as for the non-hierarchical queuing and are defined in Table 8-23. The maximum allowable burst for hierarchical egress queuing is 131040.

Copyright© 2022 Ciena® Corporation

With hierarchical egress queuing, the EIR of an element is able to inherit the EIR of the next higher level element. For example, if the default queue group EIR policy is set to inherit, then the EIR of the default queue group is equal to the shaper rate of the port. Similarly, if a default queue group queue's EIR policy is set to inherit, then the EIR of that queue is equal to the default queue group EIR. When the EIR changes for an element, any element below whose EIR policy is set to inherit also gets that new EIR value. If the EIR policy is set to configure, then the configured EIR of that element cannot exceed the EIR of the next higher level element. This prevents the over-subscription of the EIR. For example, if the default queue's EIR policy is set to configure, then the configured EIR for that default queue cannot exceed the operational EIR of the default queue group.

The sum of the CIRs cannot exceed the CIR or shaper rate of the next higher level element. For example, the sum of the CIRs for the default queues of a port cannot exceed the CIR of the default queue group. Similarly, the sum of the CIRs for all the queue groups (default and custom) of a port cannot exceed the shaper rate of the port.

The treatment of rates for aggregation is similar to non-hierarchical egress queuing.

WRED can only be applied to the queues so the congestion avoidance profile is only set for each custom or default queue.

Configuration objects

Hierarchical egress queueing comprises these configuration objects:

- Egress-queue-group profile
- Dynamic-egress-queue-group profile
- Egress-queue-group
- Egress-port

The egress-queue-group-profile represents the characteristics of the queue group. The profile can be associated with multiple queue groups. Any change in the profile is propagated to the associated queue group. There is a default profile automatically created called “Default-Profile”. These attributes are configurable for the queues:

- CIR
- EIR
- CBS
- EBS
- Scheduler weight
- Congestion avoidance profile

Copyright© 2022 Ciena® Corporation

- EIR policy (inherit or configure)

The dynamic-egress-queue-group-profile allows the network operator to specify the per-queue CIR and EIR attributes of the dynamic egress queue group profile in percentage values relative to the absolute values of the CIR and EIR attributes in the parent custom egress queue group port instance. 39XX/51XX switches automatically adjust the actual per-queue CIR and EIR values at runtime without the need to change the corresponding dynamic egress queue group profile's per-queue CIR and EIR configuration.

There is a default egress-queue-group profile named “Default-Profile” that is automatically created. This profile has queue attributes of EIR policy set to inherit, CIR=0, CBS=0, EBS=256 and EIR=1000000 for each queue. The congestion avoidance profile is set to the default congestion avoidance profile named Default-S-WRED. The scheduler weights of the queues are as follows:

- Queue 0:20
- Queue 1:30
- Queue 2:40
- Queue 3:50
- Queue 4:60
- Queue 5:70
- Queue 6:80
- Queue 7:90

Each egress-queue-group is associated with an egress-queue-group-profile. An egress-queue-group can be attached to multiple ports. When a port is added to an egress-queue-group, these attributes can be configured:

- Priority
- Scheduler weight
- CIR
- EIR
- CBS
- EBS
- EIR policy (inherit or configure)

For each port, there is an egress-port. The egress-port represents the port scheduler. The egress-port object contains the attributes of the non-hierarchical egress-port-queue-group object. The attributes of the egress-port include:

- Scheduler algorithm

Copyright© 2022 Ciena® Corporation

- Shaper rate
- Burst size
- Traffic shaping (on/off)
- Dynamic max queue depth

The default queue group is represented by the egress-port-queue-group object. the attributes for the egress-port-queue-group are different for hierarchical egress queueing versus non-hierarchical egress queuing. The hierarchical attributes are:

- CIR
- EIR
- CBS
- EBS
- Priority
- Scheduler weight
- EIR policy

> ***Note:*** When configuring service based on queuing, certain configuration changes may cause traffic disruption on the port. This occurs when associating or disassociating a queue map or queue group to a VS/VLAN, changing priority of a queue group, changing scheduler algorithm, adding/removing a port from a queue group.

Local CPU-generated control, management and OAM traffic QoS treatment

This feature allows network operators to tune the RCoS to egress queue and the RCoS to FCoS mapping for locally generated control, management and OAM frames. Different network operators provide varying QoS treatment to control, management and OAM PDUs in their networks. They require a control to tune the RCoS to egress queue mapping table for local CPU-generated frames.

Up to now, the egress CoS queue for frames generated by the local CPU has been hard-coded to 7. Therefore, all the local CPU-generated frames egress from CoS-7, no matter the RCoS of the frame.

> ***Note:*** This feature is only applicable to frames generated by the local CPU and *not* for frames generated by FPGA.

Copyright© 2022 Ciena® Corporation

Figure 8-14 shows how a network operator can assign different egress queues for control, management and OAM PDUs. Another network operator can give the same QoS treatment, such as CoS-7, to all frames generated by the local CPU.

Figure 8-14
Different egress queues for control, management, and OAM PDUs

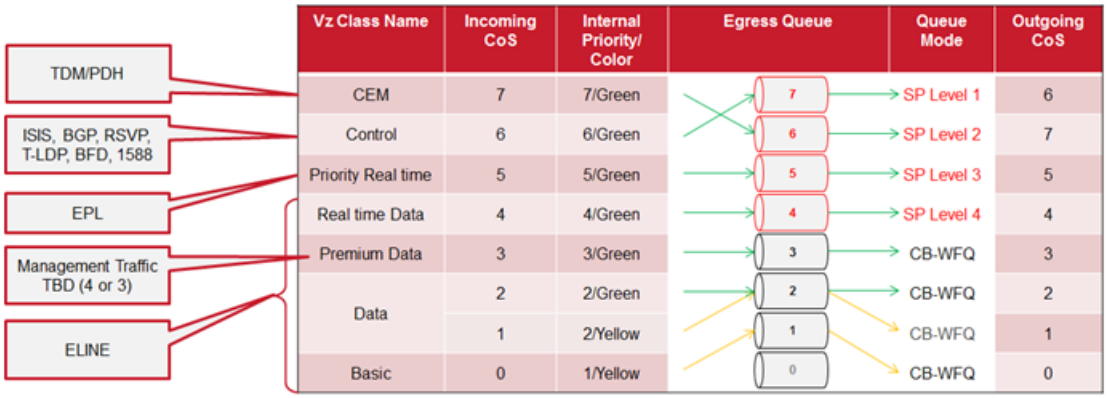

DSCP and PCP assignment per protocol

Control, management and OAM PDUs generated by the local CPU are assigned a default DSCP or PCP value using the protocol-cos mapping table, if this mapping table is enabled. MPLS TC for MPLS frames is always equal to the frame PCP value.

A protocol to CoS mapping table is provided to configure the CoS value per protocol. The protocol to CoS mapping table is global to the device and by default, this table is configured with default values. The mapping table can be enabled or disabled and if it is not enabled, then the device uses the existing behavior to set the DSCP/PCP values for different protocols. For protocols with only l2 headers, a CoS value in the range 0-7 is used, whereas protocols with l3 headers use a CoS value in the range 0-63.

PCP of the l3-protocol (control, management and OAM) is either derived from DSCP (mapping table) or by using existing behavior. A control is provided to users to choose how the PCP of l3 (control, management and OAM) frames can be derived. L3 protocol PCP settings from DSCP mapping can only be set when the protocol-cos mapping table is enabled. As per the current design, PCP of the l3-protocol frames generated by the local CPU is derived from the interface priority, which is configurable per interface.

Copyright© 2022 Ciena® Corporation

Inner tag and outer tag priority of the double-tagged frame is modified using DSCP mapping when the protocol-cos-map is enabled and PCP control is set to “dscp-mapping”. For the ARP protocol, inner tag and outer tag PCP are derived from the protocol-cos-map table when this map is enabled.

This table captures DSCP/PCP values per protocol when the protocol-cos-map table is enabled or disabled.

Table 8-25
DSCP and PCP per protocol behavior on initiator nodes

Protocol	Mapping table disabled			Mapping table enabled				
	DSCP (IPv4/ IPv6)	PCP	MPLS-TC	DSCP IPv4/ IPv6 (Default values)	PCP setting - DSCP mapping		PCP setting - Interface-pri	
					PCP (Default)	MPLS-TC	PCP	MPLS-TC
ARP	N/A	interface-pri (7)	N/A	N/A	6	N/A	6	N/A
BGP (Keepalives, TCP Ack)	48	interface-pri (7)	N/A	48	6	N/A	interface-pri (7)	N/A
IS-IS	N/A	7	N/A	N/A	6	N/A	6	N/A
LDP (Keepalives, Hello, TCP Ack)	48	interface-pri (7)	N/A	48	6	N/A	interface-pri (7)	N/A
RSVP (PATH, RESV)	0	interface-pri (7)	N/A	48	6	N/A	interface-pri (7)	N/A
PTP (1588)	configurable	configurable	N/A	48	6	N/A	configurable	N/A
Telnet (TCP SYNC, ACK, DATA)	0	interface-pri (7)	N/A	24	3	N/A	interface-pri (7)	N/A
DHCP (L3 relay)	0	interface-pri (7)	N/A	24	3	N/A	interface-pri (7)	N/A
SSH	0	interface-pri (7)	N/A	24	3	N/A	interface-pri (7)	N/A
SNMP TRAP	0	interface-pri (7)	N/A	24	3	N/A	interface-pri (7)	N/A
Syslog	0	interface-pri (7)	N/A	24	3	N/A	interface-pri (7)	N/A
DHCP-client	0	interface-pri (7)	N/A	24	3	N/A	interface-pri (7)	N/A

Copyright© 2022 Ciena® Corporation

Table 8-25
DSCP and PCP per protocol behavior on initiator nodes

Protocol	Mapping table disabled			Mapping table enabled				
	DSCP (IPv4/ IPv6)	PCP	MPLS-TC	DSCP IPv4/ IPv6 (Default values)	PCP setting - DSCP mapping		PCP setting - Interface-pri	
					PCP (Default)	MPLS-TC	PCP	MPLS-TC
Single Hop BFD	56	7	Not supported	48	6	Not supported	6	Not supported
Multi-Hop BFD	0	interface-pri (7)	Not supported	48	6	Not supported	interface-pri (7)	Not supported
ICMP	0	interface-pri (7)	N/A	24	3	N/A	interface-pri (7)	N/A
IP Traceroute	0	interface-pri (7)	N/A	24	3	N/A	interface-pri (7)	N/A
MPLS Ping (GAL, IP)	0	7	7	48	6	6	6	6
MPLS VC Ping (GAL, IP)	0	7	7	48	6	6	6	6

Note 1: PTP PCP and DCSP configuration are overwritten by the protocol-rcos mapping table values if this table is enabled.

Note 2: The PCP of DHCP-L2-Relay frames is not modified using the protocol-cos-map table.

Determining RCoS from DSCP or PCP

If the protocol-cos-map is enabled, then the RCoS value is derived from the frame DSCP or PCP value. If the Layer 3 (l3) header is present in the frame, then RCoS is derived from the DSCP value, whereas if the l3 header is not present, then RCoS is derived from the outer tag PCP value. If the protocol-cos-map is disabled, then RCoS is derived as per existing behavior, that is, RCoS is derived from L2 PCP and not from DSCP. In most cases, the default PCP is set to 7 and therefore, RCoS is derived as 7.

For local CPU-generated MPLS frames, MPLS-TC is always equal to the outer tag PCP value and PCP for MPLS OAM frames is always derived from DSCP if the protocol-cos-map is enabled.

FCoS to RCoS mapping is always one-to-one for CPU-generated frames (the protocol list is not limited to the protocol-cos-map table, refer to all frames generated by the CPU) and is *not* configurable when the protocol-cos-map is

Copyright© 2022 Ciena® Corporation

enabled (even when modifying DefaultFcosRcos map). When the protocol-cos-map is disabled, then RCoS is derived from the DefaultFcosRcos map for CPU-generated frames.

Table 8-26
CPU port FCoS RCoS table

L3 DSCP	L2 PCP	MPLS TC	RCoS
0-7	0	0	0
8-15	1	1	1
16-23	2	2	2
24-31	3	3	3
32-39	4	4	4
40-47	5	5	5
48-55	6	6	6
56-63	7	7	7

RCoS to egress queue mapping

RCoS to egress queue assignment is configurable per device but *not* configurable per protocol level. All protocols with the same RCoS are given the same QoS treatment. The CPU port RCoS map is created at initialization time and all RCoS (0-7) is mapped to egress queue-7. Users can modify the RCoS to queue mapping table using the management interface.

The CPU port RCoS queue map cannot be associated with any physical/aggregation port and cannot be used for the custom queue rcos-queue map. Also, the CPU port RCoS queue map cannot be deleted or renamed.

The default CPU port RCoS map is:

```
+-------------------------- RCOS TO QUEUE MAPPING ----------------------------+
|                                                                              |
+-------------------------------+----------------------------------------------+
| Name                          | CpuPort-RCOS                                 |
| Id                            | 1                                            |
| Type                          | RCOS To Queue Mapping                        |
+------------------------------------------------------------------------------+
| RCOS:    | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |                                  |
|          +---+---+---+---+---+---+---+---+                                   |
| Queue:   | 7 | 7 | 7 | 7 | 7 | 7 | 7 | 7 |                                  |
+----------+---+---+---+---+---+---+---+---+-------------------------------------+
```


Copyright© 2022 Ciena® Corporation

RCoS to FCoS mapping at egress port

The current behavior of SAOS 6.x is supported for the egress-rcos-fcos map.

Remote management interface CoS behavior

Protocol-cos-map table settings take precedence over remote management interface QoS settings (priority and DSCP) when the mapping table is enabled. When the protocol-cos-map table is disabled, then remote management QoS settings take precedence.

When the protocol-cos-map is enabled and the PCP setting is configured as “interface-priority”, then the PCP of the frame is derived from the user-configured priority of the remote management interface.

Copyright© 2022 Ciena® Corporation

This table captures the behavior of DSCP and PCP for the remote management interface:

Table 8-27
Remote management interface DSCP/PCP behavior for initiator and responder nodes

Protocol	Mapping table disabled			Mapping table enabled				
	DSCP (IPv4/ IPv6)	PCP	MPLS-TC	DSCP (IPv4/ IPv6) (Default values)	PCP setting - DSCP mapping		PCP setting - Interface-pri	
					PCP	MPLS-TC	PCP	MPLS-TC
ARP	N/A	remote-interface-pri (7)	PCP (7)	N/A	6	equal to PCP	6	equal to PCP
Telnet (TCP SYNC, ACK, DATA)	interface-default-dscp (0)	remote-interface-pri (7)	PCP (7)	24	3	3	remote-interface-pri (7)	PCP (7)
SSH	interface-default-dscp (0)	remote-interface-pri (7)	PCP (7)	24	3	3	remote-interface-pri (7)	PCP (7)
SNMP TRAP	interface-default-dscp (0)	remote-interface-pri (7)	PCP (7)	24	3	3	remote-interface-pri (7)	PCP (7)
Syslog	interface-default-dscp (0)	remote-interface-pri (7)	PCP (7)	24	3	3	remote-interface-pri (7)	PCP (7)
DHCP-client	interface-default-dscp (0)	remote-interface-pri (7)	PCP (7)	24	3	3	remote-interface-pri (7)	PCP (7)
ICMP	echo-reply (IPv4): preserve echo-reply (IPv6): interface-default-dscp (0) other types: interface-default-dscp (0)	remote-interface-pri (7)	PCP (7)	24	3	3	remote-interface-pri (7)	PCP (7)
IP Traceroute	IPv4: 48 IPv6: 0	remote-interface-pri (7)	PCP (7)	24	3	3	remote-interface-pri (7)	PCP (7)

Transit node behavior for DSCP and PCP

If a packet is copied to the CPU on the transit node for some special packet processing, then the DSCP of the frame is preserved and is *not* modified by

Copyright© 2022 Ciena® Corporation

the protocol-cos map table. The RSVP and DHCP L3 relay protocols do special handling of frames on the transit node.

To provide more flexibility, three different modes are defined for transit nodes:

- Mode 1: protocol-cos-map is disabled—Current behavior of DSCP/PCP.
- Mode 2: protocol-cos-map is enabled—Use the DSCP from the protocol-cos-map table.
- Mode 3: protocol-cos-map is enabled—Preserve the DSCP of the ingress frame.

The user can select Mode 2 or Mode 3 using the management interface. By default, the DSCP mode is set to “map” and can be modified to “preserve”. The DSCP mode is only applicable when protocol-cos-map is enabled. When the DSCP mode is set to “preserve”, then only DSCP is preserved and PCP is derived either from dscp-mapping or interface-priority as mentioned in Table 8-28.

These modes are defined only for the protocols listed in this table.

Table 8-28
DSCP and PCP per protocol behavior on transit nodes

	Mode 1		**Mode 2 (map)**			**Mode 3 (preserve)**		
Protocol	**DSCP/PCP (protocol-cos-map disabled)**		**DSCP/PCP (protocol-cos-map enabled)**			**DSCP/PCP (protocol-cos-map enabled)**		
	DSCP	**PCP**	**DSCP (Default values)**	**PCP—DSCP mapping**	**PCP—interface pri**	**DSCP**	**PCP—DSCP mapping**	**PCP—interface pri**
RSVP	0	interface-pri (7)	48	6	interface-pri (7)	preserve	From DSCP6	interface-pri (7)
DHCP L3 relay	0	interface-pri (7)	24	3	interface-pri (7)	preserve	From DSCP6	interface-pri (7)

RCoS is not derived from the ingress-port-rcos map. Instead, RCoS and egress queue behavior is the same as defined in “Determining RCoS from DSCP or PCP” on page 8-64 and “RCoS to egress queue mapping” on page 8-65.

Note: Preservation mode is *not* supported for the RSVP protocol.

Request or reply type frames DSCP and PCP behavior

This section describes the DSCP and PCP behavior to be expected with request or reply type frames.

Copyright© 2022 Ciena® Corporation

DSCP behavior of the reply type frame

DSCP of the reply type frames is preserved and is *not* modified by the protocol-cos map table. Due to SAOS limitations, DSCP of the reply type frames is assigned from the protocol-cos-map table if this table is enabled. RCoS and egress queue of the reply type frame is recalculated as described in "Determining RCoS from DSCP or PCP" on page 8-64 and "RCoS to egress queue mapping" on page 8-65. RCoS of reply type frames is not determined based on the FCoS to RCoS table on the ingress port.

> ***Note:*** Due to SAOS limitations, DSCP of the reply frame cannot be preserved for some of the protocols. See Table 8-29 for details.

Information on which protocols support or do not support preservation of DSCP are found in this table:

Table 8-29
DSCP behavior on responder nodes

Protocol	Request/Reply type messages supported	DSCP preservation supported	DSCP of Reply frame (IPv4/IPv6) (protocol-cos-map disabled)	DSCP of Reply frame (protocol-cos-map enabled, default values)
ARP	Yes	N/A	N/A	N/A
BGP (Keepalives, TCP Ack)	Yes	No	48	From protocol-cos-map (48)
IS-IS	No	N/A	N/A	N/A
LDP (Keepalives, Hello, TCP Ack)	Yes	No	48	From protocol-cos-map (48)
RSVP (PATH, RESV)	Yes	No	0	From protocol-cos-map (48)
Telnet (TCP SYNC, ACK, DATA)	Yes	No	0	From protocol-cos-map (24)
DHCP (L3 Relay)	Yes	No	DHCP reply (DSCP depends on DHCP server configuration/ implementation)	DHCP reply (DSCP depends on DHCP server configuration/ implementation)
SSH	Yes	No	0	From protocol-cos-map (24)
SNMP TRAP	No	N/A	N/A	N/A
Syslog	No	N/A	N/A	N/A
DHCP-client	No	N/A	N/A	N/A
Single Hop BFD	No	N/A	N/A	N/A

Copyright© 2022 Ciena® Corporation

Table 8-29
DSCP behavior on responder nodes

Protocol	Request/Reply type messages supported	DSCP preservation supported	DSCP of Reply frame (IPv4/IPv6) (protocol-cos-map disabled)	DSCP of Reply frame (protocol-cos-map enabled, default values)
Multi-Hop BFD	No	N/A	N/A	N/A
ICMP	Yes	No	echo-reply (IPv4): preserve echo-reply (IPv6): 0 other types: 0	From protocol-cos-map (24)
IP Traceroute (reply frame is ICMP protocol)	Yes	No	IPv4: 48 IPv6: 0	From protocol-cos-map for ICMP protocol (24)
MPLS Ping (GAL, IP)	Yes	Yes	preserve	preserve
MPLS VC Ping (GAL, IP)	Yes	Yes	preserve	preserve

The DSCP mode for reply frames cannot be modified by the user and default values are shown in Table 8-29.

PCP behavior of the reply type frame

PCP of the reply type frames is derived from interface-priority if the protocol-cos-map is disabled. If the protocol-cos-map is enabled on the responder node, then PCP behavior is the same as described in "DSCP and PCP assignment per protocol" on page 8-62.

Copyright© 2022 Ciena® Corporation

For MPLS Ping and MPLS VC Ping protocols, PCP is preserved for reply type frames when the path is set to LSP. When the path is set to IPv4, then behavior for PCP preservation for reply type frames is like other protocols, that is, from interface priority no matter the protocol-cos-map operating status, whereas MPLS-TC is preserved.

Table 8-30
PCP behavior on responder nodes

Protocol	Mapping table disabled		Mapping table enabled			
	PCP	MPLS-TC	PCP setting - DSCP mapping		PCP setting - Interface-pri	
			PCP (Default)	MPLS-TC	PCP	MPLS-TC
ARP	interface-pri (7)	equal to pcp	6	equal to pcp	6	equal to pcp
LDP (Keepalives, Hello, TCP Ack)	interface-pri (7)	n/a	6	n/a	interface-pri (7)	n/a
RSVP (PATH, RESV)	interface-pri (7)	n/a	6	n/a	interface-pri (7)	n/a
PTP (1588)	configurable	n/a	6	n/a	configurable	n/a
Telnet (TCP SYNC, ACK, DATA)	interface-pri (7)	n/a	3	n/a	interface-pri (7)	n/a
DHCP (L3 Relay)	interface-pri (7)	n/a	3	n/a	interface-pri (7)	n/a
SSH	interface-pri (7)	n/a	3	n/a	interface-pri (7)	n/a
ICMP	interface-pri (7)	n/a	3	n/a	interface-pri (7)	n/a
IP Traceroute	interface-pri (7)	n/a	3	n/a	interface-pri (7)	n/a
MPLS Ping (GAL, IP) – LSP Path	preserve	preserve	preserve	preserve	preserve	preserve
MPLS Ping (GAL, IP) – IPV4 Path	interface-pri (7)	preserve	interface-pri (7)	n/a	interface-pri (7)	n/a
MPLS VC Ping (GAL, IP) – LSP Path	preserve	preserve	preserve	preserve	preserve	preserve
MPLS VC Ping (GAL, IP) – IPV4 Path	interface-pri (7)	preserve	interface-pri (7)	n/a	interface-pri (7)	n/a

Copyright© 2022 Ciena® Corporation

RCoS and egress queue assignment for untagged traffic

CPU-generated untagged frames can be categorized in any the following:

- Protocol defined in the protocol-cos-map table
- Protocol not defined in the protocol-cos-map table

If the protocol is defined in the protocol-cos-map, then the RCoS value of untagged frames is configured using the protocol-cos-map. If the protocol is not defined in the mapping table, then the RCoS value is 7 and cannot be configured.

The egress queue is derived from the RCoS to queue mapping table for the CPU port. Even if the protocol is not defined in the protocol-cos-map table, then the egress queue is assigned using the CPU port queue-map table corresponding to r-cos-7 (RCoS is fixed but the queue can be configured).

Configuration example

See "Sample configuration for local CPU-generated control, management and OAM traffic QoS treatment" on page 8-258 for details.

Limitations

- All the transit and local CPU-generated protocols use the egress queue from the RCoS to queue mapping table configured for the CPU port, even if protocols are not part of protocol-rcos mapping table. For example, CCM, LM/DM, R-APS, L2CFT, and so on.
- The CPU port RCoS queue map is overridden by the custom queue rcos-queue map if any custom queue is attached to the service and local CPU generates any traffic for that service.The only exception is frames that are directly injected on the egress port (bypassing the ingress and egress pipeline of the device).
- Single-hop-BFD and multi-hop-BFD are subdivisions of IP-BFD and do not cover LSP BFD and VCC BFD.
- DSCP of the reply frame cannot be preserved for some of the protocols and is assigned from the protocol-cos-map table when this mapping table is enabled. See Table 8-28 on page 8-68 for details.
- The SSH server can use any user-defined port number for listening in SAOS. But this feature is supported only for standard SSH port 22. In this case, all the response packets from the SSH server to the SSH client have a DSCP bit valued 0.
- The MPLS-TC of CPU-generated frames is not modified by the RCoS to FCoS table.
- Management traffic from the local interface (out of band management) is also stamped with DSCP from the protocol-cos-map when this map is enabled.

Copyright© 2022 Ciena® Corporation

- The PCP of all the protocols (with L3 headers) generated by the local CPU port are derived from DSCP when the PCP control is set to “dscp-mapping” (protocol list is not restricted by protocols governed by the protocol-cos-map table).

Dynamic microflow detection and greedy flow monitoring

A flow can be considered as a traffic stream that is identified based on certain fields (like Src/Dst Mac or Src/Dst IP and so on). The traffic-profiling feature (as supported in SAOS 6.17.1 and later) provides options to meter predefined flows, that is, the user needs to configure flow qualifiers and then devices meter the same flow. There are cases when flows might not be known in advance and need to be identified and detected as and when present in the network. Once a flow is identified and detected, then it can be metered and policed. Due to a limited number of resources in 6.x-based devices, it is not possible to apply metering for all detected flows. Instead, metering can be applied to certain bandwidth-hungry flows, called greedy flows.

The user needs to define qualifiers for the parent flow, which may have multiple microflows. Microflows need not be predefined and are detected based on certain fixed attributes (Src/Dst MAC, Src/Dst IPv4 Address, Src/Dst IPv6 Address, MPLS Labels, IPv4 over MPLS, and so on) of incoming traffic. Once a microflow is detected, then it can be associated as a child profile with a parent profile for monitoring and policing purposes. It may be required that newly detected microflows under a parent flow are to be metered at certain predefined rates. But it may not be possible to meter all the detected microflows due to hardware resource constraints. In this case, some of the newly detected microflows can be monitored for a certain time period to check their bandwidth usage. If these flows are consuming bandwidth over a certain threshold limit, then these flows can be called greedy or bandwidth-hungry flows. These greedy flows can be metered till their bandwidth utilization becomes normal.

The main purpose of this feature is to identify greedy microflows and police them until their bandwidth consumption is under a given threshold.

This figure shows a possible use case:

Copyright© 2022 Ciena® Corporation

Figure 8-15
Possible use case

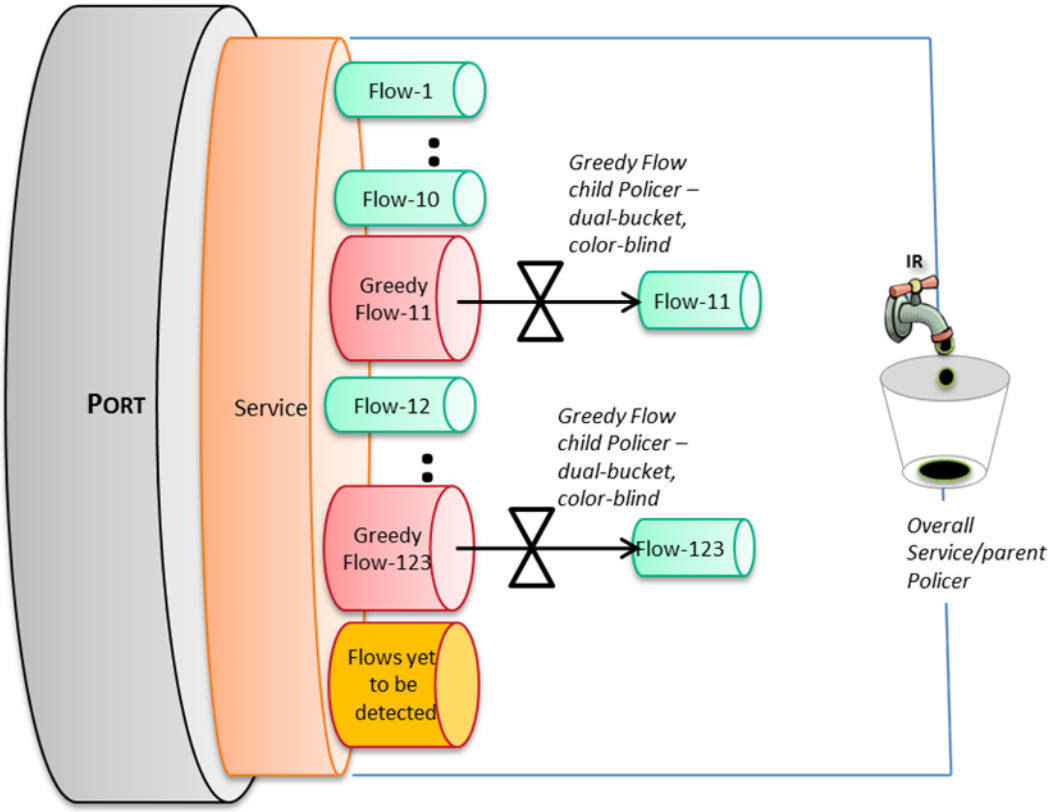

The following points help to understand the Possible use case figure:

- The EVC/service is configured on the port.
- The service is being metered and policed by the parent meter.
- An unknown number of microflows can be part of a given service.
- A certain number of microflows are detected by the device.
- Some of the microflows are identified as greedy flows and are being metered by child policers.
- Some of the microflows are yet to be detected by the device. These flows may also have some greedy flows, which will be identified with continuous operation of this feature.

This figure provides an introduction to greedy flow monitoring operations:

Copyright© 2022 Ciena® Corporation

Figure 8-16
Overview of how greedy flow monitoring works

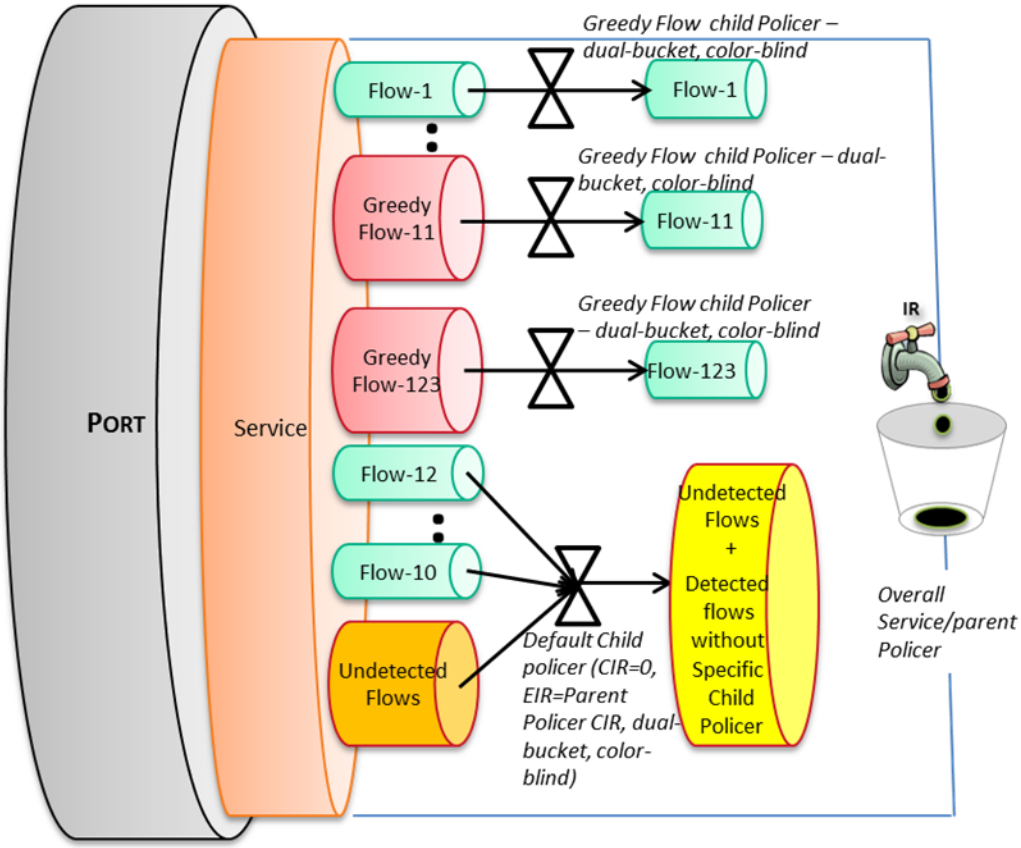

The following points help to understand the Overview of how greedy flow monitoring works figure:

- The EVC/service is configured on the port.
- The service is being metered and policed by the parent meter.
- An unknown number of microflows can be part of a given service.
- A certain number of microflows are detected by the device.
- Some of the microflows are being monitored to check if they are normal flows or greedy flows. For example, flow 1.
- Some of the microflows are identified as greedy flows and are being metered by child policers.
- Some of the microflows are yet to be detected by the device. These flows may also have some greedy flows, which will be identified with continuous operation of this feature.

Copyright© 2022 Ciena® Corporation

- A default child policer is used for all flows except greedy flows and flows under monitoring. This child policer is needed to detect new flows and does not rate-limit the traffic.

More information about how greedy flow monitoring works is provided in the Basic working section.

Basic working

Dynamic microflow detection and greedy flow monitoring is dependent on traffic-profiling and uses the same resources to identify and monitor parent/macroflows and child/microflows.

This feature can be divided into the following parts for understanding purposes:

- Monitoring of parent/macro/service flow bandwidth utilization to enable and disable flow detection.
- Detection and identification of new microflows.
- Monitoring the detected microflows to identify greedy microflows.
- Metering the greedy microflows.
- Removing non-greedy microflows from the database.

This flow diagram explains the high-level working of steps mentioned in forthcoming sections:

Copyright© 2022 Ciena® Corporation

Figure 8-17
High-level working steps

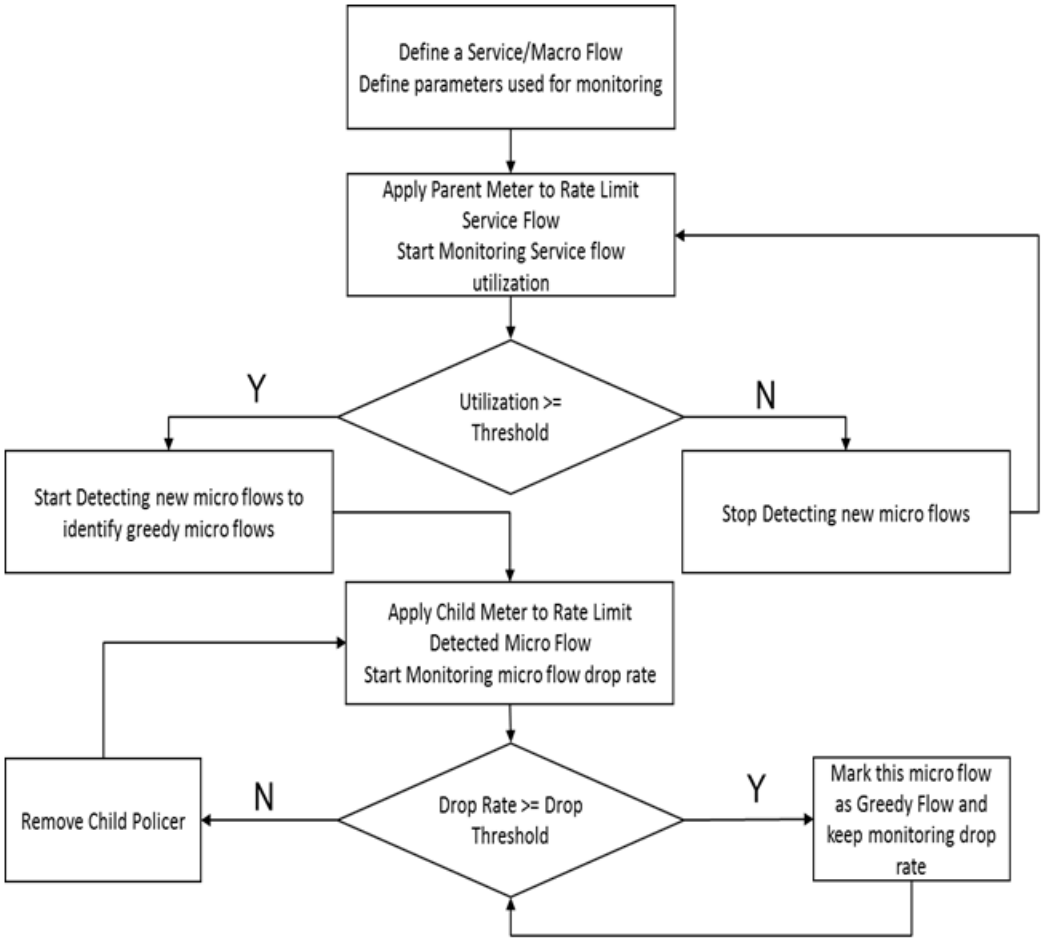

This flow diagram shows the high-level stage transitions related to a microflow:

Copyright© 2022 Ciena® Corporation

Figure 8-18
High-level stage transitions

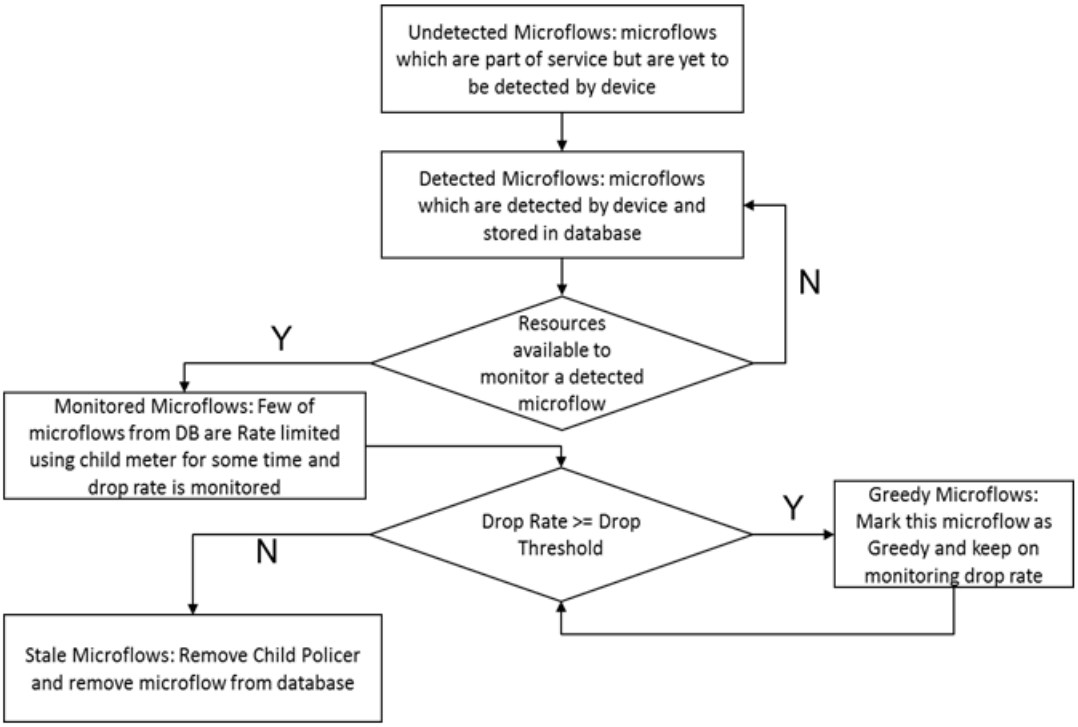

Monitoring of service/parent macroflow

A service/parent or macroflow can be considered as an envelope that can contain multiple child/microflows. For example, a port, port + VLAN, and port + VS-based flow can have multiple microflows. A parent traffic profile is used to rate-limit and monitor service flow. The statistics of the parent traffic profile are checked at regular intervals to calculate utilization:

Utilization (Bits per second) =
*(((Accepted Bytes at T2 – Accepted Bytes at T1) * 8) / (T2-T1))*

Detection of microflows only occurs during the period when utilization is over threshold. This helps to reduce CPU load. Monitoring of detected flows and metering of greedy flows does not stop even if utilization is below threshold and cools off gradually with time or as greedy flows become normal flows.

Detection and identification of microflows

Flows that are part of the service and can be identified by certain fields in frames are called microflows. A device can identify a microflow by parsing additional fields from the L2 and L3 headers of a frame. The following fields can be used to identify a microflow:

- Source MAC address

Copyright© 2022 Ciena® Corporation

- Destination MAC address
- Source IP address
- Destination IP address
- MPLS labels
- IPv4 fields (Src IP and Dst IP) in MPLS payload

Note: MPLS frames are only parsed if the MPLS label stack is maximum three labels deep, that is, the bottom-of-the-stack (BoS) label is present within the first three labels. If the MPLS labels stack is more than three labels deep, then only the Src and Dst MAC addresses are used for identifying the microflow.

This figure shows the packet fields (in yellow) that need to be parsed for identifying new microflows:

Figure 8-19
Packet fields parsed for identifying new microflows

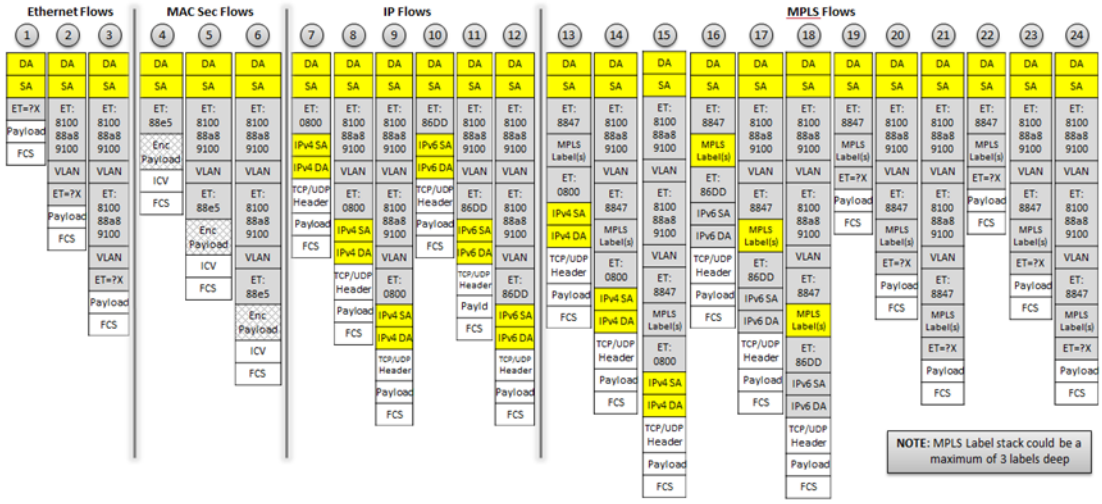

Packet field parsing and overlapping microflows

Dynamic microflow detection and greedy flow monitoring parses fields from incoming traffic to identify a flow. Packet/frame parsing is not done for frames where fields of frames are not valid. For example, for IPv4 frames: IP version, IP header length, Src/Dst IP must be valid. The same is true for IPv4 payload in IPv4 over MPLS cases. For IPv6 frames: IP version and Src/Dst IP must be valid. For all flows where the frame's fields are not valid, Src/Dst MAC addresses and Ether Type are used to identify and rate-limit the microflows.

Copyright© 2022 Ciena® Corporation

Since microflow detection is dynamic in nature, there can be overlapping cases, too. Consider a Non-IPv4 over MPLS case. This flow is identified and rate-limited using Src/Dst MAC and Ether Type. Now if another IPv4 over MPLS flow with the same Src/Dst MAC starts, then it is not detected as a separate flow and is rate-limited by the first flow. Note that if IPv4 over MPLS flow is detected first, then IPv6 over MPLS flow is also detected as a separate flow. This happens due to classifier overlap. There can be other such scenarios, as well. Basically, if a Src/Dst MAC address and Ether Type (IPv4, IPv6 and MPLS) based flow is detected first, then other flows with the same parameters but different L3, L4 or MPLS fields are not detected separately and are rate-limited by the same flow that was detected first.

Monitoring the detected microflow to identify greedy flows

Once a microflow is detected by the device, a child traffic profile is used to meter the flow. Microflow is metered with the same traffic profile parameters as needed for greedy flow. Statistics of the child traffic profile are checked on regular intervals to calculate the drop rate:

Drop rate (Bits per second) =
*(((Dropped Bytes at T2 – Dropped Bytes at T1) * 8) / (T2-T1))*

A microflow can be marked as greedy flow if the drop rate remains above a given threshold for a configurable time. In case the drop rate remains below threshold for a configurable time, then it is considered as normal, non-greedy flow and the child traffic profile is removed for this flow. All the detected microflows are periodically monitored to identify greedy microflows.

Metering of greedy microflows

A microflow that is identified as greedy microflow is rate-limited by the same child traffic profile as used to monitor the same microflow. A greedy flow is rate-limited until its drop rate remains above threshold. In case its drop rate remains less than threshold for the configurable time, then it is no longer rate-limited by the child traffic profile. This flow remains in the database and can be monitored again.

Removing non-greedy microflows from the database

At a given point, a device can have the following categories of microflows under a service/parent flow:

- Undetected microflows.
- Microflows that are detected and stored in the database.
- Microflows that are detected and are being policed/metered (monitored) to check if some of these can be potential greedy microflows or can be marked as normal microflows.
- Microflows that are identified as greedy microflow and are being rate-limited.

Copyright© 2022 Ciena® Corporation

- Normal microflows that are waiting to be flushed out of the database.

Due to the limited number of resources, only a limited number of microflows can be kept in the database and can be monitored over a regular interval. But a service can have many microflows, so to detect greedy microflows efficiently, microflows are regularly moved in and out of the database. All the microflows that are monitored and marked normal are removed from the database to make room for new microflows.

General feature interactions

This section discusses how dynamic microflow detection and greedy flow monitoring interact with other SAOS 6.x general features.

Traffic-profiling

Dynamic microflow detection and greedy flow monitoring use the same resources as used by traffic-profiling. Standard profile, envelope, and global profile cannot be configured on the same device as a dynamic traffic profile is configured.

Resource-manager

As mentioned in Traffic-profiling, this feature uses traffic-profiling resources, but these resources must be allocated prior to configuring this feature. The user needs to reboot the device in case resources are modified when the profile-mode is set to global-l2-l3-extended. Setting the profile-mode to global-l2-l3-extended is not allowed in case resources are modified (to a non-zero value) without rebooting the device. Ciena recommends the user to reserve twice the increment step size of the traffic-profiling resource bucket. (For example, if the increment step size of traffic profiling is 512, then the user ideally must reserve 2 x 512=1024 traffic-profiling resources to use this feature.)

Packet capture

The rate at which traffic samples can be given to the host CPU is very limited and is shared by packet capture and this feature. The rate that can be configured is in the range of 1-N packets per second. N is platform-dependent.

Aggregation port

This feature can be used by aggregation ports to detect microflows for traffic ingressing on aggregation ports. Member ports can be added/removed in aggregation through configuration or member ports can become distributing and non-distributing based on the Link Aggregation Control Protocol (LACP) state. All the microflows that are detected are not flushed immediately and are aged out after microflow monitoring cycle completion.

QinQ VS

All microflows related to a QinQ VS are immediately flushed in case the S-VLAN of QinQ VS gets modified due to a configuration change.

Copyright© 2022 Ciena® Corporation

Limitations

- All the limitations of global traffic-profiling are also applicable to dynamic microflow detection and greedy flow monitoring.
- This feature cannot be enabled on ports that are also used by the benchmark feature.
- This feature cannot be enabled on ports that are also used by the packet capture feature.
- This feature cannot be enabled on TDM ports.
- Since detection of microflows and identification of greedy flow is software-based, this feature can impact the host CPU performance.
- The host CPU can process only a limited number of traffic samples, thus microflow detection is not deterministic and is statistical.
- Identification of a greedy microflow may take some time due to the limited number of resources.
- Even though hardware resources are available, a device cannot police new greedy microflows if the sum of CIRs of all existing greedy microflows and microflows under monitoring exceeds the total CIR of service flows. Since CIR values are normalized before they are applied in hardware, there can be cases where “service-flow-cir” and “greedy-microflow-cir” are configured in such a way that the total number of microflows that can be rate-limited becomes one less than the (service-flow-cir/greedy-microflow-cir) count.
- A device does not validate the checksum of L3 and L4 fields of a sample frame that is used to detect new microflows.

Recommended ports

Although dynamic microflow detection and greedy flow monitoring can be enabled on any given port (including aggregation ports), Ciena recommends this feature be used with UNI and ENNI ports only. Use of this feature with NNI ports and MPLS terminating ports is *not* recommended.

Configuration

This section discusses items relevant to the configuration of dynamic microflow detection and greedy flow monitoring.

Copyright© 2022 Ciena® Corporation

Configurable parameters and usage

Dynamic microflow detection and greedy flow monitoring allows the user to create dynamic profiles and associate service flow with a profile. Some global parameters and some profile-specific parameters are explained in this table:

Table 8-31
Configurable parameters and their usage

Global configurable parameters			
Parameter	**Default value**	**Range**	**Usage**
service-flow-monitor-timer	1 second	1-10 seconds	The service/parent flow statistics are monitored at the interval given by this timer value. Utilization is calculated and compared against the service-flow-utilization-threshold.
microflow-monitor-timer	3 seconds	1-10 seconds	Micro/child flow statistics are monitored at the interval given by this timer value. The drop rate is calculated and compared against the greedy-microflow-drop-threshold.
greedy-microflow-monitor-timer	5 seconds	1-10 seconds	The greedy micro/child flow statistics are monitored at the interval given by this timer value. The drop rate is calculated and compared against the greedy-microflow-drop-threshold.
microflow-batch-monitor-count	10	1-20	The maximum number of detected microflows that can be monitored simultaneously using child meters to identify greedy flows is controlled by this parameter.
Per profile configurable parameters			
Parameter	**Default value**	**Range**	**Usage**
service-flow-cir	100 Mbps	Not applicable.	CIR (Kbps) to be used for parent meter.
service-flow-cbs	256 KB	Not applicable.	CBS (KB) to be used for parent meter.
service-flow-utilization-threshold	60 Mbps	Not applicable.	Bandwidth utilization (Kbps) of service/parent flow is compared to this threshold (Kbps) to check if the service flow bandwidth is overutilized and greedy flow detection needs to be enabled.
service-flow-over-threshold-count	5	1-100	Service flow utilization must remain above the service-flow-utilization-threshold for these many times, continuously, to mark that the flow is overutlizing allocated bandwidth.
service-flow-under-threshold-count	5	1-100	Service flow utilization must remain below the service-flow-utilization-threshold for these many times, continuously, to clear the overutilization situation.

Copyright© 2022 Ciena® Corporation

Table 8-31
Configurable parameters and their usage (continued)

Per profile configurable parameters (continued)			
greedy-microflow-cir	10 Mbps	-	CIR (Kbps) to be used for the child meter.
greedy-microflow-cbs	256 KB	-	CBS (KB) to be used for the child meter.
greedy-microflow-drop-threshold	1 Mbps	-	The drop rate (Kbps) of the child/microflow is compared against this threshold (Kbps) to check if the microflow is greedy or not.
greedy-microflow-under-threshold-count	5	1-100	The drop rate of the greedy microflow must remain below the greedy-microflow-drop-threshold for these many times, continuously, to mark the flow as normal flow.
microflow-over-threshold-count	5	1-100	The drop rate of the monitored microflow must remain above the greedy-microflow-drop-threshold for these many times, continuously, to mark the microflow as greedy microflow.
microflow-under-threshold-count	5	1-100	The drop rate of the monitored microflow must remain below the greedy-microflow-drop-threshold for these many times, continuously, to mark the microflow as non-greedy microflow.
microflow-monitor-count	5	1-1000	A detected microflow is monitored these many times before getting cleaned up from the microflow database.

Associating a service flow with the dynamic profile

To associate a service flow with a profile, a port (or range of ports), virtual switch and C-VLAN need to be added in the profile. The port is a mandatory parameter, whereas VLAN and VS are optional. A port or range of ports is considered as a single association and a port can be added in multiple profiles, given that there is at least one non-overlapping qualifier in all associations.

Examples of non-overlapping associations include:

- Port1 and VS1 can be added in Profile1.
- Port1 and VS2 can be added in Profile2.
- Port1 and VS3 can be added in Profile1.
- Port1 not allowed.
- Port1 and VLAN1 not allowed.

Copyright© 2022 Ciena® Corporation

- Port1 and VS4 and VLAN1 not allowed.

All the associations related to a port must have the same classifiers. For example, in case the first association is added with Port and VS, then all associations to this port need to have VS only. Any association like Port, VLAN or Port are not allowed and are considered overlapping.

> ***Note:*** Both VS and VLAN cannot be added in a single association and are mutually exclusive classifiers.

Copyright© 2022 Ciena® Corporation

Devices supported

The SAOS 6.x devices supported for dynamic microflow detection and greedy flow monitoring are listed in this table:

Table 8-32
SAOS 6.x devices supported for dynamic microflow detection and greedy flow monitoring

Resource	Number	SAOS 6.x Device
A device supports a maximum of N service/macroflows.	8	5142, 5160
	4	3924, 3926, 3928, 3942
	2	3903, 3904, 3905
A device stops detection of microflows after detecting N microflows.	2048	5142, 5160
	1024	3942
	512	3903, 3904, 3905, 3924, 3926, 3928
A device starts detection of new microflows if: **1** Detected microflows are less than N and **2** Service flow utilization is more than the configured threshold.	2048	5142, 5160
	1024	3942
	512	3903, 3904, 3905
A device only supports traffic-profiling classifier resource allocation in multiples of N through the resource-manager when the profile-mode is set to global-l2-l3-extended.	1024	5142, 5160
	512	3903, 3904, 3905, 3924, 3926, 3928, 3942
A device only supports profile-mode as global-l2-l3-extended when traffic-profiling classifier resources are allocated in multiples of N through the resource-manager.	1024	5142, 5160
	512	3903, 3904, 3905, 3924, 3926, 3928, 3942
A device supports N classifier resources for dynamic-policing.	Where N is the number of resources allocated for traffic-profiling using the resource-manager.	All
A device reserves N classifier resources for internal use.	128	5142, 5160
	64	3924, 3926, 3928, 3942
	32	3903, 3904, 3905
A device shall consume N classifier resources to police/meter a single greedy microflow.	4	All
A device consumes N classifier resources to simultaneously monitor detected microflows to identify greedy microflows.	Where N is: The number of simultaneous microflows that need to be monitored * 4	All

Copyright© 2022 Ciena® Corporation

Table 8-32
SAOS 6.x devices supported for dynamic microflow detection and greedy flow monitoring

Resource	Number	SAOS 6.x Device
A device supports the metering of maximum N greedy microflows.	Where N is: (The number of classifier resources allocated for traffic-profiling using resource-manager – The number of classifier resources consumed for internal use – The number of classifier resources used to monitor the detected micro flows simultaneously) / 4	All
A device supports N profiles for dynamic microflow policing.	8	5142, 5160
	4	3924, 3926, 3928, 3942
	2	3903, 3904, 3905
A device supports N non-overlapping macroflow's qualifier association with a single dynamic profile or across multiple dynamic profiles.	8	5142, 5160
	4	3924, 3926, 3928, 3942
	2	3903, 3904, 3905

Procedures

Procedures related to dynamic microflow detection and greedy flow monitoring include:

- "Configuring a device for dynamic microflow detection and greedy flow monitoring" on page 8-260
- "Minimally configuring microflow policing on port 1" on page 8-263
- "Working examples of dynamic microflow detection and greedy flow monitoring" on page 8-264

Copyright© 2022 Ciena® Corporation

Procedure 8-1
Configuring Class of Services (CoS) policies on a port

You can configure CoS policies for each port.

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 RCoS and RColor values (rcos-to-l2-outer-pcp-map).

Step	Action
1	Set the resolved CoS policy: `port set port <PortNameList> [resolved-cos-policy <dot1d-tag1-cos|dot1d-tag-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© 2022 Ciena® Corporation

Example

This example sets a resolved CoS policy.

```
port set port 1 resolved-cos-policy dot1d-tag1-cos
```

This 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
```

This example sets an egress-frame-cos-policy.

```
port set port 1 egress-frame-cos-policy rcos-to-l2-outer-pcp-map
```

In this example, the outgoing frame PCP/DEI gets remarked based on the ingress port-1 FCoS -> RCoS map and egress port-2 RCos -> FCoS map.

```
port set port 1 resolved-cos-remark-l2 true
port set port 1 resolved-cos-policy dot1d-tag1-cos
port set port 2 egress-frame-cos-policy rcos-to-l2-outer-pcp-map
```


Copyright© 2022 Ciena® Corporation

Procedure 8-2
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 port <PortName> is the port that you want to display settings for.

—end—

Example

This 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© 2022 Ciena® Corporation

Procedure 8-3
Creating a resolved CoS map

This table shows the default resolved CoS map (DefaultFcosRcos).

Table 8-33
Default resolved CoS map

Frame PCP/L2 CoS 802.1D	DEI/CFI	Frame L3 DSCP CoS	RCoS	RCOLOR
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© 2022 Ciena® Corporation

You can configure 1 custom resolved CoS map for each port, including these types of mapping:

- L2 CoS - maps the customer PCP/L2 CoS 802.1D and DEI/CFI priority to RCoS and RColor.
- L3 CoS - maps the customer L3 CoS DSCP to RCoS and RColor.

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

This 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© 2022 Ciena® Corporation

Procedure 8-4
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 <NUMBER: 0..7>] {r-cos <NUMBER:0..7>} [r-color <green|yellow>]` where cos-map <String[15]> — is the identifier for the custom resolved CoS map. dot1dpri-cos <0..7> — ingress .1D value to map. dot1dpri-dei <NUMBER: 0..1> — is the ingress frame color cfi/dei value. l3-cos-dscp <NUMBER: 0..63> — is the ingress dscp value to map. [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

This 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© 2022 Ciena® Corporation

Procedure 8-5
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 RCoS and RColor 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

This example sets the resolved CoS map for a port.

```
port set port 1 resolved-cos-map custom-1 resolved-cos-remark-l2 true
```


Copyright© 2022 Ciena® Corporation

Procedure 8-6
Deleting a custom RCoS map

You can delete custom resolved CoS maps, but not the default resolved CoS map.

Note: A custom resolved CoS queue map cannot be deleted if it is assigned to a port.

Step	Action
1	Set the port using the custom resolved CoS map to the default RCoS map: `port set port <PortName> resolved-cos-map DefaultFcosRcos`
2	Delete a resolved CoS map: `traffic-services cos-mapping resolved-cos-map delete cos-map <Resolved CoS Map Name>` where cos-map <Resolved CoS Map Name> is the resolved CoS map to delete.

—end—

Example

This example deletes the custom resolved CoS map name custom-1.

```
port set port 9 resolved-cos-map Default-RCOS
traffic-services cos-mapping resolved-cos-map delete cos-map custom-1
```


Copyright© 2022 Ciena® Corporation

Procedure 8-7
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

This 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  |
| 1             |     1   ||                             ||      | green  |
|               |         || 8-15                         ||      | green  |
```


Copyright© 2022 Ciena® Corporation

```
+--------------+---------++-------------------------------+++------+--------+
|              |         ||                               ||      |        |
| 2            |    1    ||                               ||    2 | green  |
|              |         || 16-23                          ||      | green  |
...
+--------------+---------++-------------------------------+++------+--------+
|              |         ||                               ||      |        |
| 7            |    0    ||                               ||    7 | green  |
| 7            |    1    ||                               ||      | green  |
|              |         || 56-63                          ||      | green  |
+--------------+---------++-------------------------------+++------+--------+
```

This 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© 2022 Ciena® Corporation

Procedure 8-8
Creating a frame CoS map

With frame CoS mapping, the RCoS and RColor associated with a frame is mapped to specific CoS values in the frame based on the default Frame CoS mapping table (DefaultRcosFcos) as shown in this table.

Table 8-34
Default Frame CoS mapping table

RCoS	RColor	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 “global” type and 12 “l2-only” type custom frame CoS maps to use the frame’s RCoS and RColor to map the customer PCP/L2 CoS 802.1D and Discard Eligibility Indicator (DEI) priority values on egress.

Copyright© 2022 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]>] type l2-only` where cos-map <String[15]> is the identifier for the custom frame CoS map.

—end—

Example

This example creates a custom frame CoS map.

```
traffic-services cos-mapping frame-cos-map create cos-map custom-1 type l2-only
```


Copyright© 2022 Ciena® Corporation

Procedure 8-9
Modifying a frame CoS map

You can modify the attributes for the default frame CoS map 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

This 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© 2022 Ciena® Corporation

Procedure 8-10
Deleting a frame CoS map

You can delete custom frame CoS maps, except for the default frame CoS map. A custom frame CoS map that is in use cannot be deleted.

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 RCoS -> FCoS name.

—end—

Example

This example deletes a frame CoS map.

```
traffic-services cos-mapping frame-cos-map delete cos-map custom-1
```


Copyright© 2022 Ciena® Corporation

Procedure 8-11
Enabling and disabling the remarking of CPU generated frames

SAOS devices support the configuration option to control PCP/DEI remarking of CPU generated frames. By default, the system does not perform PCP/DEI remarking of CPU generated frames based on the egress port RCoS to RCoS map.

Step	Action
1	Enable the remarking of CPU generated frames: `traffic-services cos-mapping set cpu-cos-remark-l2 <on|off>` where cpu-cos-remark-l2 <on|off> — enables or disables the remarking of CPU generated frames.

—end—

Example

This example enables the remarking of CPU generated frames.

```
traffic-services cos-mapping set cpu-cos-remark-l2 on
```


Copyright© 2022 Ciena® Corporation

Procedure 8-12
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 RCoS 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. When egress Frame CoS mapping is enabled, the ingress setting resolved-cos-remark-l2 is overridden and no longer preserves 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

This 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© 2022 Ciena® Corporation

Procedure 8-13
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

This 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                          |
| type              |  global                     |
+-------+---------+++---------+---------+---------+
|  RCOS | RCOLOR  |||  L2-COS |  L2-DEI | MPLS-TC |
+-------+---------+++---------+---------+---------+
|     0 |  green  |||     0   |     0   |     0   |
|     0 |  yellow |||     0   |     0   |     0   |
|     1 |  green  |||     1   |     0   |     1   |
|     1 |  yellow |||     1   |     0   |     1   |
|     2 |  green  |||     2   |     0   |     2   |
|     2 |  yellow |||     2   |     0   |     2   |
|     3 |  green  |||     3   |     0   |     3   |
|     3 |  yellow |||     3   |     0   |     3   |
|     4 |  green  |||     4   |     0   |     4   |
|     4 |  yellow |||     4   |     0   |     4   |
|     5 |  green  |||     5   |     0   |     5   |
|     5 |  yellow |||     5   |     0   |     5   |
|     6 |  green  |||     6   |     0   |     6   |
|     6 |  yellow |||     6   |     0   |     6   |
|     7 |  green  |||     7   |     0   |     7   |
|     7 |  yellow |||     7   |     0   |     7   |
+-------+---------+++---------+---------+---------+
```


Copyright© 2022 Ciena® Corporation

Procedure 8-14
Configuring ingress RCoS to egress queue mapping

Each physical port supports 8 CoS queues. Depending on the policies applied to the port, the RCoS and RColor values may change during processing. The final internal RCoS value is used to map the frame to one of the 8 physical port CoS queues as shown in this table.

Table 8-35
Default internal RCoS mapping

RCoS	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 3 custom Queue Map Profiles map RCoS to CoS queue. The mapping-mode of custom queue map profiles is always RCoS mapped. To configure a port to use a custom queue map profile, set the *ingress-to-egress-qmap* attribute. RCoS values (0 - 3) can be mapped to CoS Queues (0 - 3), or any combination thereof (for example, all 3 RCoS values mapped to the same CoS queue).

Step	Action

To create and modify a custom RCoS to CoS queue table and assign it to an egress port

1 Create a custom RCoS mapping for a queue:

```
traffic-services queueing queue-map create rcos-map
<rcos-map>>
```

where

rcos-map <rcos-map>	is the RCoS queue map to be created.

Copyright© 2022 Ciena® Corporation

2 Set an RCoS mapping for a queue:

```
traffic-services queueing queue-map set rcos-map <rcos-map> {rcos <NUMBER>} {queue <NUMBER: 0-7>}
```

where

rcos-map <NUMBER>	is the RCoS range value.
queue <NUMBER: 0-7>	is the queue value.

3 Assign the queue to an egress port:

```
port set port <port> ingress-to-egress-qmap <RCoS Queue Map>
```

where

port <port>	is the port to set.
ingress-to-egress-qmap<RCoS Queue Map>	sets the ingress-to-egress queue map.

To set the RCoS map to the default for a port

4 Set the RCoS map to the default for a port:

```
port set <port> ingress-to-egress-qmap <Default-RCoS>
```

where

port <port>	is the port to set.
ingress-to-egress-qmap<Default-RCos>	is the default RCoS profile.

Note: You cannot unset a port's ingress-to-egress-qmap profile, as this is a required parameter. To remove a custom RCoS profile from a port you must replace it with another RCoS profile (for example, the Default-RCOS profile).

—end—

Example

This example creates and modifies a custom RCoS to CoS queue table and assigns it to an egress port.

```
traffic-services queuing queue-map create rcos-map custom-map-1
traffic-services queueing queue-map set rcos-map custom-map-1 rcos 0 queue 3
port set port 9 ingress-to-egress-qmap custom-map-1
```

This example set the RCoS map to the default for a port:

```
port set port 9 ingress-to-egress-qmap Default-RCOS
```


Copyright© 2022 Ciena® Corporation

Procedure 8-15
Displaying an RCoS map

Display an RCoS map.

Step	Action
1	To display an RCoS map: `traffic-services queuing queue-map show [rcos-map <RcosQueueMap>]` where rcos-map <RcosQueueMap> is the RCoS map to be displayed. **—end—**

Example

This example shows output for the default RCoS 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© 2022 Ciena® Corporation

Procedure 8-16
Applying RCoS 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 RCoS as shown in this table.

Table 8-36
Custom map

Frame PCP/ L2 CoS 802.1D	Frame DEI/ CFI	RCoS	RColor
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 are assigned an RCoS 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 RCoS 1.

Copyright© 2022 Ciena® Corporation

Figure 8-20
RCos policies and mapping

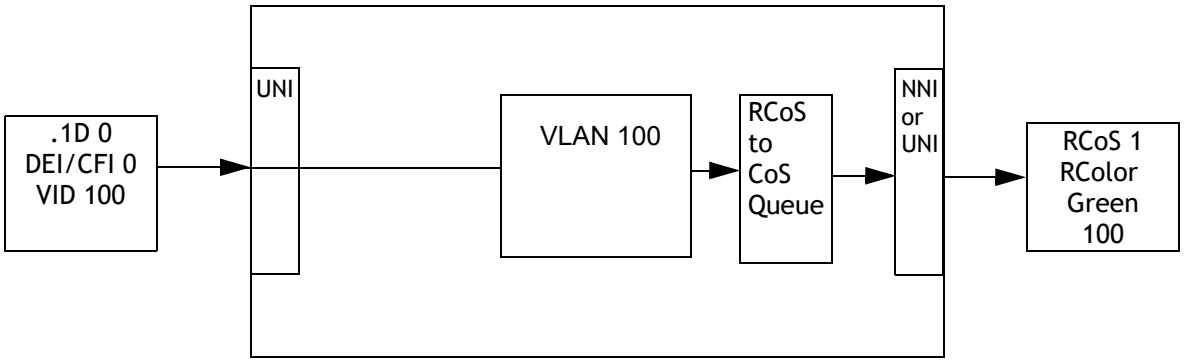

Step	Action
1	Create the RCoS map and set the mapping for PCP/L2 CoS.1D to RCoS values. `traffic-services cos-mapping resolved-cos-map create cos-map <cos-map>` where cos-map <cos-map> is the FCoS -> RCoS map name.
2	Set the FCoS to RCoS map entry: `traffic-services cos-mapping resolved-cos-map set cos-map <cos-map> [dot1dpri-cos <dot1d values and ranges>] {r-cos <NUMBER: 0-7>} [r-color <green | yellow>]` where cos-map <cos-map> sets the FCoS to RCos map entry. dot1dpri-cos <dot1d values and ranges> is the ingress.1D value to map. r-cos <NUMBER: 0-7> is the RCoS value to use. r-color <green | yellow> is the resolved color.

Copyright© 2022 Ciena® Corporation

3 Configure the ingress port's resolved CoS policy and map:

```
port set port <port> [resolved-cos-policy <dot1dtag1-cos|dot1d-tag2-cos|fixed-cos|l3-dscp-cos>] [resolved-cos-map <FCos to RCos map>]
```

where

port <port>	is the port to set.
resolved-cos-policy <dot1dtag1-cos|dot1d-tag2-cos|fixed-cos|l3-dscp-cos>>	is the resolved CoS policy.
resolved-cos-map <FCos to RCos map>	selects the resolved CoS map.

4 Create the provider VLANs and add the NNI port:

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

where

vlan <vlan>	is the VLAN being created.
name <String>	is the name of the VLAN.

5 Add the VLAN to the port:

```
vlan add vlan <vlan> {port <Port list>}
```

where

vlan <vlan>	is the VLAN being created.
port <Port list>	is the list of ports to add to the VLAN.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example creates the RCoS map and sets the mapping for PCP/L2 CoS.1D to RCoS 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
```

This example configures the ingress port's resolved CoS policy and map.

```
port set port 2 resolved-cos-policy dot1d-tag1-cos resolved-cos-map MyL2Map
```

This example creates the provider VLANs and adds the NNI port.

```
vlan create vlan 100,1000-1001
vlan add vlan 1000-1001 port 6
```


Copyright© 2022 Ciena® Corporation

Procedure 8-17
Applying RCoS policies and mapping in a virtual switch

Apply RCoS policies and map in a virtual switch.

Step	Action
1	Create the RCoS map: `traffic-services cos-mapping resolved-cos-map create cos-map <cos-map>`
2	Set the mapping for DSCP to RCoS values: `traffic-services cos-mapping resolved-cos-map set cos-map <cos-map> [l3-cos-dscp <dscp values and ranges>] [dot1dpri-cos <dot1d values and ranges] [dot1dpri-dei <NUMBER: 0-1>] [mpls-tc <mpls-tc values and ranges>]{r-cos <NUMBER: 0-7>} [r-color <green|yellow>]` where cos-map <cos-map> — is the FCoS to RCoS name. l3-cos-dscp <dscp values and ranges> — is the ingress DSCP value to map. dot1dpri-cos <dot1d values and ranges — is the ingress.1D value to map. dot1dpri-dei <NUMBER: 0-1> — is the ingress frame color cfi/dei value. mpls-tc <mpls-tc values and ranges> — is the mpls tc/exp value to map. r-cos <NUMBER: 0-7> — is the RCoS value to use. r-color <greenlyellow> — is the resolved color.
3	Configure the resolved CoS Policy and map of the ingress port: `port set port <port> resolved-cos-policy {<dot1d-tag1-cos|fixed-cos|l3-dscp-cos>} {resolved-cos-map <FCOS to RCOS Map>}`
4	Create the provider VLANs: `vlan create vlan <vlan>`

Copyright© 2022 Ciena® Corporation

5 Add the NNI port:

```
vlan add vlan <vlan> port <port>
```

6 Create the virtual circuits.

```
virtual-circuit ethernet create vc <vc name> {vlan
<VLAN>} [statistics <on|off>]
```

7 Reserve VLANs for the virtual switches.

```
virtual-switch add reserved-vlan <VLAN ID list>
```

```
virtual-switch ethernet create vs <vs> vc <vc>
```

8 Add the members with the encap policy to inherit from the port settings.

```
virtual-switch ethernet add vs <vs> port <Port list>
encap-cos-policy port-inherit
```

9 Set the untagged data virtual switch for the member port.

```
port set port <port> untagged-data-vs <Virtual Switch
Name>
```

—end—

Copyright© 2022 Ciena® Corporation

Example

With the example in this section, frames with CVID 100 and specific DSCP values are assigned an internal RCoS as shown in the Custom map table, because the virtual switch member inherits the custom RCoS map of the port. Untagged frames are assigned an RCoS according to the port's fixed RCoS configuration.

Table 8-37
Custom map

Frame DSCP	RCoS	RColor
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 RCoS 7 and RColor Yellow.

Figure 8-21
Frame with CVID 100 and DSCP value of 0 mapped to RCoS 7 and RColor Yellow

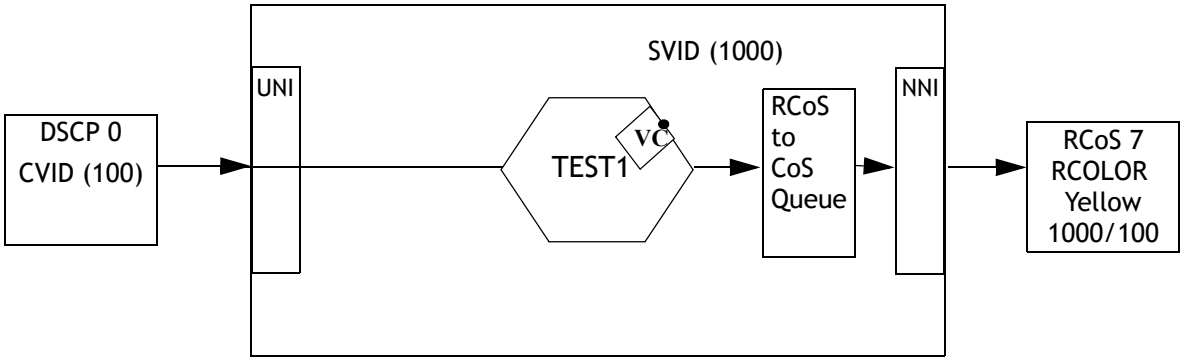

Copyright© 2022 Ciena® Corporation

Procedure 8-18
Configuring an L2-only type RCoS to FCoS map

Configure an L2-only type RCoS to FCoS map.

Note: If the type is not specified, the default type global is created.

Step	Action
1	Configure an L2-only type RCoS to FCoS map: `traffic-services cos-mappng frame-cos-map create cos-map <name> type l2-only` where cos-map <name> is the cos-map name.
2	Set RCoS and RColor to L2-COS 5 and L2-DEI: `traffic-services cos-mapping frame-cos-map set cos-map <RCoS to FCoS map> dot1dpri-cos <NUMBER: 0..7> r-color <green | yellow> dot1dpri-dei <NUMBER: 0..1> rcos <NUMBER: 0..7>` where cos-map <RCoS to FCoS map> specifies the RCoS -> FCoS name. dot1dpri-cos <NUMBER: 0..7> specifies the egress frame CoS (pcp) value. r-color <greenlyellow> specifies the resolved color. dot1dpri-dei <NUMBER: 0..1> specifies the egress frame color. rcos <NUMBER: 0..7> is the RCoS value.
3	Attach custom map to a port: `port set port <port> frame-cos-map`

—end—

Example

This example shows how to configure an L2-only type RCoS to FCoS map.

```
> traffic-services cos-mapping frame-cos-map create c1 type l2-only
> traffic-services cos-mapping frame-cos-map set cos-map c1 r-cos 3 dot1dpri-
cos 5 r-color green dot1dpri-dei 0
> port set port 5 frame-cos-map c1
```


Copyright© 2022 Ciena® Corporation

Procedure 8-19
Configuring a global type RCoS to FCoS map

1 Configure a global type RCoS to FCoS map:

```
traffic-services cos-mapping frame-cos-map create cos-map
<name>
```

where

cos-map <name> is the cos-map name.

2 Set RCoS and RColor to L2-COS and L2-DEI:

```
traffic-services cos-mapping frame-cos-map set r-cos
<cos-map <RCoS to FCoS map> dot1dpri-cos <NUMBER: 0..7>
r-color <green | yellow> dot1dpri-dei <NUMBER: 0..1> r-
cos <NUMBER: 0..7> mpls-tc <NUMBER: 0..7>
```

where

cos-map <RCoS to FCoS map>	specifies the RCoS -> FCoS name.
dot1dpri-cos <NUMBER: 0..7>	specifies the egress frame CoS (pcp) value.
r-color <green|yellow>	specifies the resolved color.
dot1dpri-dei <NUMBER: 0..1>	specifies the egress frame color.
r-cos <NUMBER: 0..7>	is the RCoS value.
mpls-tc <NUMBER: 0..7>	is the mpls-tc value.

3 Attach custom map to a port:

```
port set port <port> frame-cos-map
```

—end—

Example

This example shows how to configure a global type RCoS to FCoS map.

```
> traffic-services cos-mapping frame-cos-map create c2
> traffic-services cos-mapping frame-cos-map set cos-map c2 r-cos 2 dot1dpri-
cos 6 r-color green dot1dpri-dei 0 mpls-tc 4
> port set port 5 frame-cos-map c2
```


Copyright© 2022 Ciena® Corporation

Procedure 8-20
Displaying the RCoS to FCoS map

Display the RCoS to FCoS map.

Step	Action
1	Display the RCoS to FCoS map: `traffic-services cos-mapping frame-cos-map show cos-map <name>` where cos-map <name> is the cos-map name.

—end—

Example

This example shows the output for the traffic-services cos-mapping frame-cos-map show command for the cos-map C1.

```
> traffic-services cos-mapping frame-cos-map show cos-map c1

+---------------------------------------+
|   RESOLVED COS MAP TO FRAME COS INFO  |
| Name                |  c1             |
| logical-id          |  2              |
| type                |  l2-only        |
|  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  |||    5    |    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© 2022 Ciena® Corporation

Procedure 8-21
Enabling and disabling standard traffic profiling

By default, traffic profiling status is disabled globally and per port. To configure traffic profiling, EIR/PIR, and CIR values must be provisioned.

Step	Action

To enable traffic profiling

1 Enable traffic profiling:

```
traffic-profiling enable [<PortNameList>]
```

where

port <PortNameList> is the port that you want to enable traffic profiling on.

To disable traffic profiling

2 Disable traffic profiling:

```
traffic-profiling disable [<PortNameList>
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 8-22
Enabling and disabling standard traffic extended resources

By default, standard traffic extended resources are disabled globally and per port.

If extended resources is globally disabled, “extended” resource-type traffic-profiles cannot be created. If “extended” resource-type traffic-profile is already created, then feature cannot be disabled until all the “extended” resource-type traffic-profiles are deleted.

While creating the traffic-profile, resource-type can be optionally provided but if the type is not provided and feature is enabled then Software will determine the type of the traffic-profile in such a manner that maximum number of traffic-profiles can be created.

Step	Action
To enable extended resources	
1	Enable extended resources: `traffic-profiling enable extended-resource`
To disable extended resources	
2	Disable extended resources: `traffic-profiling disable extended-resource`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 8-23
Configuring an envelope

Configuring an envelope consists of configuring the envelope on a port, and configuring the bandwidth-flow-coupling flag and the bandwidth-flow-statistics.

Note 1: Envelope meters can be attached to a port in global mode. Envelope and global traffic profile cannot be supported on the same port or if the global traffic profile is created for all ports.

Note 2: The envelope is created internally on the Benchmark port under test if the envelope is attached on the same port and enabled.

Note 3: You cannot create an envelope with statistics ON if there is no resource available for counters.

Step	Action
1	Configure a traffic profiling envelope on a port: `traffic-profiling envelope create envelope <envelope-name> port <port> bw-flow-coupling-flag <on|off> bw-flow-statistics <on|off>` where envelope <envelope-name> — is the name of the envelope. port <port> — is the port number. The port must be a member of the associated virtual switch to that envelope. The port can also be an aggregation port. bw-flow-coupling flag <on|off> — sets the bandwidth-flow-coupling flag to on or off. bw-flow-statistics <on|off> — sets the bandwidth-flow statistics to on or off. The counter resource is released and statistics are cleared when statistics is set to OFF for the envelope. The default is ON.

—end—

Example

This example configures an envelope.

```
> traffic-profile envelope create envelope ENV1 port 5 bw-flow-coupling-flag
on bw-flow-statistics on
```


Copyright© 2022 Ciena® Corporation

Procedure 8-24
Configuring a bandwidth profile flow in an envelope

Configuring an envelope consists of these tasks:

- Creating an envelope
- Adding bandwidth profile flows
- Enabling the envelope

Overlapping bandwidth profile flows cannot exist in an envelope or across the envelopes attached to the UNI port. All the bandwidth flows in an envelope must satisfy the same classification criteria:

- VS
- VS + RCoS
- VS + CE-VLAN PCP

The VS is created before adding it as a classifier to the bandwidth profile blow.

The bandwidth can be shared in the envelope, but sharing of bandwidth between envelopes is not allowed.

You can add or remove a bandwidth profile flow when the envelope is already enabled. Traffic is impacted. The sharing of bandwidth may be affected. You cannot remove the last bandwidth profile flow from the envelope if the envelope is administratively enabled.

Note 1: Repeat this procedure for each bandwidth profile flow.

Note 2: The CIR max of the highest rank member in an envelope is not used.

Copyright© 2022 Ciena® Corporation

Step	Action
1	Configure a bandwidth profile flow to an envelope:

```
traffic-profiling envelope add envelope <envelope-name>
bw-profile-flow [name <flow-name>] [vs <Virtual Switch
Ethernet Name>]{r-cos <NUMEBR: 0..7> | dot1dpri <NUMBER:
0..7> | rank <NUMBER: 0..7>} {[cir <Kbps>], [eir <Kbps>],
[pir <Kbps>], [cir-max <Kbps>] [eir-max <Kbps>], [pir-max
<Kbps>]} [cbs <KBYTES>] [ebs <KBYTES>] [ingress-color-
aware <on|off>]
```

where

envelope <name>	is the name of the envelope.
bw-profile-flow [name <flow-name>]	is the bw-profile flow name.
vs <Virtual Switch Ethernet Name>	is the virtual switch name.
r-cos <NUMBER: 0..7> | dot1dpri <NUMBER: 0..7> | rank <NUMBER: 0..7>}	is the RCoS value.
cir <Kbps>	is the cir value.
eir <Kbps>	is the eir value.
pir <Kbps>	is the pir value.
cir-max <Kbps>	is the cir-max value.
eir-max <Kbps>	is the eir-max value.
pir-max <Kbps>	is the pir-max value.
cbs <Kbytes>	is the cbs value.
ebs <Kbytes>	is the ebs value
ingress-color-aware <on|off>	turns ingress-color-aware on or off.

—end—

Example

This example configures a bandwidth profile to an envelope.

```
> traffic-profile envelope add envelope ENV1 vs vs1 r-cos 0 cir 20 cir-max 50
color-aware off
```


Copyright© 2022 Ciena® Corporation

Procedure 8-25
Clearing statistics for envelope flows

Clear statistics for envelope flows.

Step	Action
1	Clear statistics for envelope flows: `traffic-profiling envelope clear statistics`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 8-26
Deleting an envelope

Delete an envelope.

Step	Action
1	Delete an envelope: `traffic-profiling envelope delete envelope <name>` **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 8-27
Enabling an envelope flow

You cannot enable the envelope if there is no port associated with the envelope. The port is mandatory to enable the envelope. The default is disabled.

You cannot enable the envelope when no bandwidth profile flow is added to that envelope.

You cannot enable the envelope if a Benchmark test is running on the same port on which the envelope is attached.

Step	Action
1	Enable an envelope flow: `traffic-profiling envelope enable envelope <name>`

—end—

Example

This example enables the envelope ENV1.

```
> traffic-profiling envelope enable envelope ENV1
```


Copyright© 2022 Ciena® Corporation

Procedure 8-28
Disabling an envelope flow

Counters of all the bandwidth profile flows in the envelope are stopped and cleared when an envelope is disabled.

When traffic-profiling is administratively disabled, envelopes are operationally disabled.

You cannot disable the envelope if a Benchmark test is running on the same port on which the envelope is attached.

Step	Action
1	Disable an envelope flow: `traffic-profiling envelope disable envelope <name>`

—end—

Example

This example disables the envelope ENV1.

```
> traffic-profiling envelope disable envelope ENV1
```


Copyright© 2022 Ciena® Corporation

Procedure 8-29
Removing an envelope flow

Remove an envelope flow.

Step	Action
1	Remove an envelope flow: `traffic-profiling envelope remove envelope <name> bw-profile-flow <flow-name>` where envelope <name> — is the envelope you are removing. bw-profile-flow <flow-name> — is the bandwidth profile-flow name.

—end—

Example

This example removes the flow BPF1 from the envelope ENV1.

```
> traffic-profiling envelope remove envelope ENV1 bw-profile-flow BPF1
```

This example removes the flow BPF2 from the envelope ENV2.

```
> traffic-profiling envelope remove envelope ENV2 bw-profile-flow BPF2
```


Copyright© 2022 Ciena® Corporation

Procedure 8-30
Modifying bandwidth flow attributes

Modify an envelope and its flow attributes:

- bandwidth flow coupling flag
- bandwidth flow statistics
- port

Note: You cannot detach a port from an envelope if the envelope is already enabled. You can modify the port when the envelope is enabled. The port must be a member of the virtual switch when you are modifying the port for an active envelope.

Step	Action
1	Modify an envelope and its flow attributes: `traffic-profiling envelope set envelope <name> bw-flow-coupling-flag <on|off> bw-flow-statistics <on|off> port <port>` where envelope <name> — is the envelope you are modifying the attributes for. bw-flow-coupling-flag <on|off> — sets the bandwidth flow coupling flag to on or off. bw-flow-statistics <on|off> — sets the bandwidth flow statistics to on or off. The default is ON. port <port> — is the port name. The port must be a member of the associated virtual switch to that envelope, and the envelope must be enabled. The port can also be an aggregation port. **—end—**

Example

This example modifies an envelope and its attributes.

```
> traffic-profiling envelope set envelope ENV1 bw-flow-coupling-flag on bw-flow-statistics off port 1
```


Copyright© 2022 Ciena® Corporation

Procedure 8-31
Displaying envelopes

You can display

- all envelopes
- envelope resource usage
- an envelope and its flow table
- statistics of envelope flows

Step	Action

To display all envelopes

1 Display all envelopes:

```
traffic-profiling envelope show
```

To display envelope resource usage

2 Display envelope resource usage:

```
traffic-profiling envelope show resources
```

To display a specific envelope

3 Display a specific envelope:

```
traffic-profiling envelope show envelope <name>
```

To display statistics of envelope flows

4 Display statistics of envelope flows:

```
traffic-profiling envelope show statistics
```

To display statistics of a specific envelope flow

5 Display statistics of a specific envelope flow:

```
traffic-profiling envelope show statistics envelope
<name>
```

Example

This example shows the output for the traffic-profiling envelope show command.

```
> traffic-profiling envelope show

+------------------------------------------------- ENVELOPES -------------------------------------+
| Envelope Name     | Admin State | Oper State | Num Of Flows | Envelope        | Ports | Stats | Flow CF |
|                   |             |            |              | Coupling Flag   |       |       | Ctrl    |
+-------------------+-------------+------------+--------------+-----------------+-------+-------+---------+
| envelope1         | Enabled     | Enabled    | 3            | Off             | 1     | On    | Off     |
+-------------------+-------------+------------+--------------+-----------------+-------+-------+---------+
| envelope2         | Enabled     | Enabled    | 3            | Off             | 1     | On    | On      |
+-------------------+-------------+------------+--------------+-----------------+-------+-------+---------+
```


Copyright© 2022 Ciena® Corporation

This example shows the output for the traffic-profiling envelope show resources command:

```
> traffc-profiling envelope show resources

+---------- ENVELOPE RESOURCE USAGE ----------+
| Parameter                        | Value      |
+----------------------------------+------------+
| Envelopes Allocated              |       4092 |
| Envelopes Used                   |          2 |
| Envelopes Available              |       4090 |
| Envelopes with Stats Allocated   |       4092 |
| Envelopes with Stats Used        |          2 |
| Envelopes with Stats Available   |       4090 |
+----------------------------------+------------+
```

This example shows the output for the traffic-profiling envelope envelope1.

```
> traffic-profiling envelope show envelope envelope1

+--------------------- ENVELOPE envelope1 ----------------------+
| Parameter                    | Value                           |
+------------------------------+---------------------------------+
| Index                        | 1                               |
| Admin State                  | Enabled                         |
| Oper State                   | Enabled                         |
| Envelope Coupling Flag       | Off                             |
| Ports                        | 1                               |
| Number Of BW Profile Flows   | 3                               |
| Stats Collection             | On                              |
| Classification Mode          | vs c-pcp                        |
+------------------------------+---------------------------------+
+------------------------------------------------- BW PROFILE FLOWS -------------------------------+
| Name      | Id | Rank | Coupling | Color |            Bandwidth(Kbps)      |    Max Burst(KB)  |  Classifiers|
|           |    |      | Flag     | Aware | CIR    | CIRmax| EIR    | EIRmax| CBS | EBS       | Type  |Value |
+----------+----+------+----------+-------+-------+-------+-------+-------+-----+-----+------+-------+------+
| envelope | 1  | 0    | Off      | On    | 20000 | 20000 | 10000 | 10000 | 256 | 128 | vs   | vs401 |      |
| 1.0      |    |      |          |       |       |       |       |       |     |     |      | c-pcp | 0    |
+----------+----+------+----------+-------+-------+-------+-------+-------+-----+-----+------+-------+------+
| envelope | 2  | 1    | Off      | On    | 15000 | 15000 | 10000 | 10000 | 256 | 128 | vs   | vs401 |      |
| 1.1      |    |      |          |       |       |       |       |       |     |     |      | c-pcp | 1    |
+----------+----+------+----------+-------+-------+-------+-------+-------+-----+-----+------+-------+------+
| envelope | 3  | 2    | Off      | On    | 15000 | 15000 | 10000 | 10000 | 256 | 128 | vs   | vs401 |      |
| 1.2      |    |      |          |       |       |       |       |       |     |     | c-pcp| 2     |      |
+----------+----+------+----------+-------+-------+-------+-------+-------+-----+-----+------+-------+------+
```


Copyright© 2022 Ciena® Corporation

This command shows the output for the traffic-profiling envelope show statistics command.

```
> traffic-profiling envelope show statistics
```

BW PROFILE FLOW STATSTICS

Envelope Name	BW Profile Flow Name	Statistics		
		Type	Bytes	Packets
envelope1	envelope1.0	Accepted	0	0
		Dropped	0	0
envelope1	envelope1.1	Accepted	0	0
		Dropped	0	0
envelope1	envelope1.2	Accepted	808044	11883
		Dropped	1231956	18117
envelope2	envelope2.5	Accepted	0	0
		Dropped	0	0
envelope2	envelope2.6	Accepted	0	0
		Dropped	0	0
envelope2	envelope2.7	Accepted	0	0
		Dropped	0	0

This command shows the output for the traffic-profiling envelope show statistics command for a specific envelope.

```
> traffic-profiling envelope show statistics envelope envelope1
```

BW PROFILE FLOW STATSTICS

Envelope Name	BW Profile Flow Name	Statistics		
		Type	Bytes	Packets
envelope1	envelope1.0	Accepted	0	0
		Dropped	0	0
envelope1	envelope1.1	Accepted	0	0
		Dropped	0	0
envelope1	envelope1.2	Accepted	808044	11883
		Dropped	1231956	18117

Copyright© 2022 Ciena® Corporation

Procedure 8-32
Setting default values for envelope and flow attributes

Set default values for envelope and flow attributes.

Step	Action

To set default values for envelope and flow attributes

1 Set default values for envelope and flow attributes:

```
traffic-profiling envelope unset envelope <envelope-name>
bw-flow-coupling-flag bw-flow-statistics port
```

where

envelope <envelope-name>	is the envelope you are setting default values for.

To set default values for bandwidth profile flow attributes

2 Set default values for bandwidth profile flow attributes:

```
traffic-profiling envelope unset envelope <envelope-name>
{bw-profile-flow <flow-name> [cir-max], [eir-max], [cbs],
[ebs], [ingress-color-aware]}
```

where

bw-profile-flow <flow-name>	is the bandwidth profile flow name you are setting default values for.

—end—

Example

This example sets the envelope ENV1 back to default values.

```
> traffic-profiling envelope unset envelope ENV1 bw-flow-coupling-flag bw-
flow statistics port
```


Copyright© 2022 Ciena® Corporation

Procedure 8-33
Configuration example for MEF 10.3 ingress metering

This example creates an envelope ENV1 with the configuration defined in "Use case 1" on page 8-35. Envelope ENV1 has 3 bandwidth profile flows for RCoS 0, 1 and 2. CE-VLAN pcp 0, 1, 2, 3 have been mapped to RCoS 0, pcp 4, 5 mapped to RCoS 1 and pcp 6, 7 mapped to RCoS 2 at UNI port 5.

Step	Action
1	Create S-VLAN: `vlan create vlan 1000`
2	Add a NNI port in S-VLAN: `vlan add vlan 1000 port 4`
3	Create a virtual circuit: `virtual-circuit ethernet create vc vc1 vlan 1000`
4	Create a virtual switch: `virtual-switch ethernet create vs vs1 vc vc1`

5	Add port 5 as an EVPL member in the virtual switch and set the encapsulation policy as port-inherit:
	`virtual-switch ethernet add vs vs1 port 5 vlan 100 encap-cos-policy port-inherit`
	`virtual-switch ethernet add vs vs1 port 5 vlan 101 encap-cos-policy port-inherit`

Copyright© 2022 Ciena® Corporation

6 Create FCoS -> RCoS map to create mapping table of C-VLAN PCP -> RCoS:

```
traffic-services cos-mapping resolved-cos map create cos-map cvlan_rcos
```

```
traffic-services cos=mapping resolved-cos-map set cos-map cvlan_rcos dot1dpri-cos 0 r-cos 0
```

```
traffic-services cos-mapping resolved-cos-map cos-map cvlan_rcos dot1dpri-cos 1 r-cos 0
```

```
traffic-services cps-mapping resolved-cos-map set cos-map cvlan_rcos dot1dpri-cos 2 r-cos 0
```

```
traffic-services cos-mapping resolved-cos-map set cos-map cvlan_rcos dot1dpri-cos 3 r-cos 0
```

```
traffic-services cos-mapping resolved-cos-map set cos-map cvlan_rcos dot1dpri-cos 4 r-cos 1
```

```
traffic-services cos-mapping resolved-cos-map set cos-map cvlan_rcos dot1dpri-cos 5 r-cos 1
```

```
traffic-services cos-mapping resolved-cos-map set cos-map cvlan_rcos dot1dpri-cos 6 r-cos 2
```

```
traffic-services cos-mapping resolved-cos-map set cos-map cvlan_rcos dot1dpri-cos 7 r-cos 2
```

7 Attach FCoS -> RCoS map to the UNI port:

```
port set port 5 resolved-cos-policy dot1d-tag1-cos
```

```
port set port 5 resolved-cos-map cvlan_rcos
```

8 Create envelope on UNI port-5:

```
traffic-profiling envelope create envelope ENV1 port 5
```

9 Add three bandwidth profile blows — H, L, M:

```
traffic-profiling envelope add envelope ENV1 vs vs1 r-cos 0 cir 20 cir-max 50 color-aware off
```

```
traffic-profiling envelope add envelope ENV1 vs vs1 r-cos 1 cir 15 cir-max 30 color-aware off
```

```
traffic-profiling envelope add envelope ENV1 vs vs1 r-cos 2 cir 15 color-aware off
```

10 Enable traffic profiling. If the traffic profiling feature is disabled, then envelope meters are operationally disabled.

```
traffic-profiling enable
```

11 Enable envelope, making it admin enable:

```
traffic-profiling envelope enable envelope ENV1
```

12 Show envelope and bandwidth profile flow:

```
traffic-profiling envelope show envelope ENV1
```


Copyright© 2022 Ciena® Corporation

Example

This example shows output from the traffic-profiling envelope show envelope ENV1 command:

```
> traffic-profiling envelope show envelope ENV1

+------------------------ ENVELOPE ENV1 ------------------------+
| Parameter                   | Value                            |
+-----------------------------+----------------------------------+
| Index                       | 1                                |
| Admin State                 | Enabled                          |
| Oper State                  | Enabled                          |
| Envelope Coupling Flag      | Off                              |
| Ports                       | 5                                |
| Number of BW Profile Flows  | 3                                |
| Stats Collection            | On                               |
| Classification Mode         | vs rcos                          |
+-----------------------------+----------------------------------+

+---------------------------------------------- BW PROFILE FLOWS ----------------------------------------+
| Name    | Id | Rank | Coupling | Color |          Bandwidth(Kbps)     |Max Burst(KB)|   Classifiers    |
|         |    |      | Flag     | Aware | CI | CIRmax | EIR  | EIRmax| CBS  | EBS  | Type  | Value    |
+---------+----+------+----------+-------+----+--------+------+-------+------+------+-------+----------+
| ENV1.0  | 1  | 0    | Off      | Off   | 20 |     50 |    0 |     0 | 128  | 128  | vs    | vs1      |
|         |    |      |          |       |    |        |      |       |      |      | rcos  | 0        |
+---------+----+------+----------+-------+----+--------+------+-------+------+------+-------+----------+
| ENV1.1  | 2  | 1    | Off      | Off   | 15 |     30 |    0 |     0 | 128  | 128  | vs    | vs1      |
|         |    |      |          |       |    |        |      |       |      |      | rcos  | 1        |
+---------+----+------+----------+-------+----+--------+------+-------+------+------+-------+----------+
| ENV1.2  | 3  | 2    | Off      | Off   | 15 |     15 |    0 |     0 | 128  | 128  | vs    | vs1      |
|         |    |      |          |       |    |        |      |       |      |      | rcos  | 2        |
+---------+----+------+----------+-------+----+--------+------+-------+------+------+-------+----------+
```


Copyright© 2022 Ciena® Corporation

Procedure 8-34
Configuration example of a c-pcp type envelope

This example creates the envelope ENV2 on port 22 which has three bandwidth profile flows for vs + c-pcp 0, 1 and 2.

Step	Action
1	Create S-VLAN: `vlan create vlan 1000`
2	Add a NNI port in S-VLAN: `vlan add vlan 1000 port 23`
3	Create a virtual circuit: `virtual-circuit ethernet create vc vc2 vlan 1000`
4	Create a virtual switch: `virtual-switch ethernet create vs vs2 vc vc2`
5	Add port 22 as an EVPL member in the virtual switch, and set the encapsulation policy as vs-inherit: `virtual-switch ethernet add vs2 port 22 vlan 200 encap-cos-policy vs-inherit`
6	Create envelope on UNI port-22: `traffic-profiling envelope create envelope ENV2 port 22`

7	Add three bandwidth profile flows:
	`traffic-profiling envelope add envelope ENV2 bw-profile-flow cir 32 cbs 4 cir-max 64 vs vs2 dot1dpri 0`
	`traffic-profiling envelope add envelope ENV2 bw-profile-flow cir 16 cbs 4 cir-max 32 vs vs2 dot1dpri 1`
	`traffic-profiling envelope add envelope ENV2 bw-profile-flow cir 16 cbs 4 cir-max 16 vs vs2 dot1dpri 2`

8	Enable traffic-profiling: `traffic-profiling envelope enable envelope ENV2`
9	Display envelope and flow parameters: `traffic-profiling envelope show envelope ENV2`

Copyright© 2022 Ciena® Corporation

Example

This example shows output from the traffic-profiling envelope show envelope EN2 command:

```
> traffic-profiling envelope show envelope ENV2
+------------------------- ENVELOPE ENV2 -------------------------+
| Parameter                    | Value                             |
+------------------------------+-----------------------------------+
| Index                        | 2                                 |
| Admin State                  | Enabled                           |
| Oper State                   | Enabled                           |
| Envelope Coupling Flag       | Off                               |
| Ports                        | 22                                |
| Number Of BW Profile Flows   | 3                                 |
| Stats Collection             | On                                |
| Classification Mode          | vs c-pcp                          |
+------------------------------+-----------------------------------+

+------------------------------------------- BW PROFILE FLOWS -------------------------------------+
| Name    | Id | Rank | Coupling | Color |          Bandwidth(Kbps)      |Max Burst(KB)|   Classifiers    |
|         |    |      | Flag     | Aware | CI | CIRmax | EIR  | EIRmax| CBS  | EBS  | Type  | Value    |
+---------+----+------+----------+-------+----+--------+------+-------+------+------+-------+----------+
| ENV2.0  | 1  | 0    | Off      | On    | 32 |     64 |    0 |     0 | 4    | 0    | vs    | vs2      |
|         |    |      |          |       |    |        |      |       |      |      | c-pcp | 0        |
+---------+----+------+----------+-------+----+--------+------+-------+------+------+-------+----------+
| ENV2.1  | 2  | 1    | Off      | On    | 16 |     32 |    0 |     0 | 4    | 0    | vs    | vs2      |
|         |    |      |          |       |    |        |      |       |      |      | c-pcp | 1        |
+---------+----+------+----------+-------+----+--------+------+-------+------+------+-------+----------+
| ENV2.2  | 3  | 2    | Off      | On    | 16 |     16 |    0 |     0 | 4    | 0    | vs    | vs2      |
|         |    |      |          |       |    |        |      |       |      |      | c-pcp | 2        |
+---------+----+------+----------+-------+----+--------+------+-------+------+------+-------+----------+
```


Copyright© 2022 Ciena® Corporation

Procedure 8-35
Configuration example of a vs-only type envelope

This example creates an envelope ENV3 on port 22 which has 3 bandwidth profile flows for 3 different VS - vs3, vs4 and vs5 with rank 0, 1 and 2.

Step	Action
1	Create S-VLAN: `vlan create vlan 1000-1005`
2	Add a NNI port in S-VLAN: `vlan add vlan 1003-1005 port 23`

3	Create a virtual circuit:
	`virtual-circuit ethernet create vc vc3 vlan 1003`
	`virtual-circuit ethernet create vc vc4 vlan 1004`
	`virtual-circuit ethernet create vc vc5 vlan 1005`

4	Create a virtual switch:
	`virtual-switch ethernet create vs vs3 vc vc3`
	`virtual-switch ethernet create vs vs4 vc vc4`
	`virtual-switch ethernet create vs vs5 vc vc5`

5	Add port 22 as an EVPL member in the virtual switch, and set the encap policy as vs-inherit:
	`virtual-switch ethernet add vs vs3 port 22 vlan 300 encap-cos-policy port-inherit`
	`virtual-switch ethernet add vs vs4 port 22 vlan 400 encap-cos-policy port-inherit`
	`virtual-switch ethernet add vs vs5 port 22 vlan 500 encap-cos-policy port-inherit`

6	Create envelope on UNI port 22: `traffic-profiling envelope create envelope ENV3 port 22`

7	Add three bandwidth profile flows:
	`traffic-profiling envelope add envelope ENV3 bw-profile-flow cir 25000 cbs 128 cir-max 25000 rank 0 vs vs3`
	`traffic-profiling envelope add envelope ENV3 bw-profile-flow cir 25000 cbs 128 cir-max 25000 rank 1 vs vs4`
	`traffic-profiling envelope add envelope ENV3 bw-profile-flow cir 25000 cbs 128 cir-max 25000 rank 2 vs vs5`

8	Enable traffic-profiling: `traffic-profiling enable`

Copyright© 2022 Ciena® Corporation

9 Display envelope and flow parameters:

```
traffic-profiling envelope show envelope ENV3
```

Example

This example shows output from the traffic-profiling envelope show envelope ENV3 command:

```
> traffic-profiling envelope show envelope ENV3

+------------------------ ENVELOPE ENV3 -------------------------+
| Parameter                    | Value                            |
+------------------------------+----------------------------------+
| Index                        | 1                                |
| Admin State                  | Enabled                          |
| Oper State                   | Enabled                          |
| Envelope Coupling Flag       | Off                              |
| Ports                        | 22                               |
| Number Of BW Profile Flows   | 3                                |
| Stats Collection             | On                               |
| Classification Mode          | vs                               |
+------------------------------+----------------------------------+
+------------------------------------------ BW PROFILE FLOWS ------------------------------+
| Name      | Id | Rank | Coupling | Color |        Bandwidth(Kbps)     |Max Burst(KB) | Classifiers    |
|           |    |      | Flag     | Aware | CIR   | CIRmax | EIR | EIRmax| CBS | EBS    | Type | Value  |
+----------+----+------+----------+-------+------+--------+-----+-------+-----+--------+-------+-------+
| ENV3.0    | 1  | 0    | Off      | On    |25000 |  25000 | 0   |   0   |  128| 0      | vs    | vs3   |
+----------+----+------+----------+-------+------+--------+-----+-------+-----+--------+-------+-------+
| ENV3.1    | 2  | 1    | Off      | On    |25000 |  25000 | 0   |   0   | 128 | 0      | vs    | vs4   |
+----------+----+------+----------+-------+------+--------+-----+-------+-----+--------+-------+-------+
| ENV3.2    | 3  | 2    | Off      | On    |25000 |  25000 | 0   |   0   | 128 | 0      | vs    | vs5   |
+----------+----+------+----------+-------+------+--------+-----+-------+-----+--------+-------+-------+
```


Copyright© 2022 Ciena® Corporation

Procedure 8-36
Enabling and disabling global profiling

Global profiles can be individually enabled or disabled.

Step	Action

To enable global profiling

1 Enable global profiling:

```
traffic-profiling global-profile enable [global-profile-
object-list]
```

where

global-profile-object-list	is the global profile you are enabling.

To disable global profiling

2 Disable global profiling:

```
traffic-profiling global-profile disable [global-
profile-object-list]
```

where

global-profile-object-list	is the global profile you are disabling.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 8-37
Enabling and disabling global parent profiling

Global parent profiles can be individually enabled or disabled.

Step	Action

To enable global parent profiling

1 Enable global parent profiling:

```
traffic-profiling global-profile parent enable [global-
profile-object]
```

where

global-profile-object	is the global parent profile you are enabling.

To disable global parent profiling

2 Enable global parent profiling:

```
traffic-profiling global-profile parent disable [global-
profile-object]
```

where

global-profile-object	is the global parent profile you are disabling.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 8-38
Setting the traffic profiling provisioning mode, profile mode, and profile type

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
```

To set the profile mode to global profile

3 Set the profile mode to global profile

```
traffic-profiling set [profile-mode <port | global-l2-l3
| global-l2-l3-extended>]
```

where

profile-mode <port | global-l2-l3 | global-l2-l3-extended>	sets the profiling mode to port, global-l2-l3, or global-l2-l3-extended.

To set the profile type for the profile mode

4 Set the profile type for the profile mode

```
traffic-profiling set [profile-type <global-profile | dynamic-
profile>]
```

where

profile-type <port | global-profile | dynamic-profile>	sets the profile type to global-profile or dynamic profile. Default is global-profile.

—end—

Copyright© 2022 Ciena® Corporation

Examples

This example sets the global profile mode.

```
traffic-profiling set profile-mode global-l2-l3-extended
```

This example sets the profile type to dynamic.

```
traffic-profiling set profile-type dynamic-profile
```


Copyright© 2022 Ciena® Corporation

Procedure 8-39
Displaying traffic profiling information in port profile mode

You can display:

- traffic profiling status
- port attributes by port
- classification mode for all ports

Step	Action

To display traffic profiling status

1 Display traffic profiling status:

```
traffic-profiling show
```

To display port attributes by port

2 Display port attributes by port:

```
traffic-profiling show port <PortName>
```

where

port <PortName> is the port that you want to display attributes for.

To display the classification mode for all ports

3 Display the classification mode for all ports:

```
traffic-profiling show classifier-mode
```

—end—

Copyright© 2022 Ciena® Corporation

Example

This example shows the 392X traffic-profiling status.

```
traffic-profiling show
+------------ TRAFFIC PROFILING GLOBAL TABLE ---------------+
| Profiling Status      | Disabled                          |
| meter-provisioning    | pir                               |
| profile-mode          | port                              |
| extended-resource     | Enabled                           |
+-------------------------------------------------------------+

+---------------- 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       |
| 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       |
+-----------+----------+----------+-------------------------+
```

This example shows the 51XX traffic-profiling status.

```
traffic-profiling show
+------------ TRAFFIC PROFILING GLOBAL TABLE ---------------+
| Profiling Status      | Disabled                          |
| meter-provisioning    | pir                               |
| profile-mode          | port                              |
+-------------------------------------------------------------+

+---------------- 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       |
| 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       |
+-----------+----------+----------+-------------------------+
```


Copyright© 2022 Ciena® Corporation

This 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                    |
+-------------------------------------------------------------+
```

This example displays the classification mode for all ports.

```
traffic-profiling show classifier-mode
+- TRAFFIC-PROFILING PORT CLASSIFIER MODE --+
| Port   |         Port Class Mode          |
+--------+----------------------------------+
| 1      |         narrow                   |
| 2      |         narrow                   |
| 3      |         narrow                   |
| 4      |         narrow                   |
| 5      |         narrow                   |
| 6      |         narrow                   |
| 7      |         narrow                   |
| 8      |         narrow                   |
| 9      |         narrow                   |
| 10     |         narrow                   |
| 11     |         narrow                   |
| 12     |         narrow                   |
+--------+----------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 8-40
Setting traffic profiling port attributes for standard profiles

Traffic profiling can be configured per port and per EVC with these 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 results in breaking an IP network. The default setting is bypass. ARP bypass is only applicable when the profiling mode is set to either standard-dscp or standard-ip-prec.

- 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 are not 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 as described in this table.

Table 8-38
Modes

Mode	Description
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. The 3903, 3903x, 3904, 3905, 3906, 3924, 3926, 3928, 3942, 5142, and 5160 platforms consume classifiers from the global meter pool.
Advanced	Supports EVC-based classification, and supports standard and hierarchical metering simultaneously. 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 on the VLAN ID.

Copyright© 2022 Ciena® Corporation

Table 8-38
Modes

Mode	Description
standard-vlan-dot1dpri	Finds a matching profiling based on the VLAN ID and outer.1D priority value in the frame.
standard-vlan-ip-prec	Finds a matching profiling based on the VLAN ID and ip-precedence value in the frame.
standard-vlan-dscp	Finds a matching profiling based on 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. Note that nonconform-standard-profile does not include untagged frames. 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, are always non-conforming because these frames do not contain the DSCP field. The default setting is drop.

  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

This 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© 2022 Ciena® Corporation

Procedure 8-41
Configuring a traffic profiling standard profile

You can perform these 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> | name <STRING>} [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|remark-green-rcos-only|remark-yellow-
rcos-only|remark-both-rcos-only>] [green-remark-rcos
<NUMBER: 0..7>] [yellow-remark-rcos <NUMBER: 0..7>]
```


Copyright© 2022 Ciena® Corporation

where

{port <port>}	is the port number.
[profile <NUMBER: 1..4096> | name <String>	is the traffic profile number or name.
[dot1dpri <NUMBER>]	is the 802.1D priority value of matching frames.
[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.

Copyright© 2022 Ciena® Corporation

where

[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 platforms operating in hierarchical-port or hierarchical-vlan mode.
[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.
[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© 2022 Ciena® Corporation

where

[remark-rcos-policy <leave|remark-green|remark-yellow|remark-both|remark-green-rcos-only|remark-yellow-rcos-only|remark-both-rcos-only>]	is the meter output color-based rcos remarking policy. Policies are: • leave: leaves the RCoS 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 RCoS value to any packets that are classified as green. The new RCoS value is assigned using green-remark-rcos <NUMBER: 0-7>. • remark-yellow: assigns a new RCoS value to any packets that are classified as yellow. The new RCoS value is assigned using yellow-remark-rcos <NUMBER: 0-7>. • remark-both: assigns a new RCoS value to any packets that are classified as green and to any packets that are classified as yellow. The new RCoS value for green packets is assigned using green-remark-rcos <NUMBER: 0-7>. The new RCoS value for yellow packets is assigned using yellow-remark-rcos <NUMBER: 0-7>. • remark-green-rcos-only: assigns a new RCoS value to any packets that are classified as green. The new RCoS value is assigned using green-remark-rcos <NUMBER:0-7>. • remark-yellow-rcos-only: assigns a new RCoS value to any packets that are classified as yellow. The new RCoS value is assigned using yellow-remark-rcos <NUMBER:0-7>. • remark-both-rcos-only: assigns a new RCoS value to any packets that are classified as green and to any packets that are classified as yellow. The new RCoS value for green packets is assigned using green-remark-rcos <NUMBER: 0-7>. The new RCoS 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.

Copyright© 2022 Ciena® Corporation

where

[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.

To modify attributes for 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 <#>], [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

This example uses EIR mode.

```
traffic-profiling set meter-provisioning eir
```


Copyright© 2022 Ciena® Corporation

```
traffic-profiling standard-profile create port 1 name TP1 cir 256 eir 512
dot1dpri 1
```

This 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
```

This 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
```

This example remarks the metered output of a traffic profile.

```
traffic-profiling standard-profile create port 1 profile 1 remark-color-
policy yellow-to-green
```

This example remarks the RCoS per traffic-profile.

```
traffic-profiling standard-profile create port 2 profile 2 remark-rcos-policy
remark-green
```

This example specifies that a traffic profile is color blind.

```
traffic-profiling standard-profile create port 1 profile 1 ingress-color-
aware off
```

This 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
```

This 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 have RCoS value of 3, so all the frames are stored in the same egress queue and at egress port they all 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 RCoS values to Yellow frames. But if the default RCoS to FCoS mapping is used, the Yellow frames in the egress port are assigned the new P-bit values and leave the port with the new P-bit value.

This example assigns all Yellow frames in a port RCoS 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© 2022 Ciena® Corporation

Procedure 8-42
Configuring a traffic profiling global traffic profile

You can perform these operations on a global traffic profile:

- create
- modify
- remove classification attributes
- delete

The CIR and PIR/EIR must be configured before you create a global traffic profile. To change profile modes, all user created traffic profiles must be deleted.

The maximum value for CIR/PIR for global profiles depends on the device (limited by the sum of the maximum bandwidth for each physical input port). See the CIR and PIR maximum values table.

Table 8-39
CIR and PIR maximum values

Device	CIR and PIR maximum values
3903	3G
3903x	3G
3904, 3905	4G
3906	3G
3924	44G
3926	62G
3928	28G
3942, 5142	60G
5160	240G

Copyright© 2022 Ciena® Corporation

Step	Action

To create a global traffic profile

1 Create a global traffic profile:

```
traffic-profiling global-profile create <global-profile-create-object> {<precedence-attr> [<dscp-attr-list>] [<dot1dpri-attr-list>] [<vlan-id-attr-list>] <cir-attr> <pir-attr> [<cbs-attr>] [<ebs-attr>] [<dscp-remark-policy-attr>] [<fixed-dscp-attr>][<untagged-attr>] [<virtual-switch-attr>] [<ingress-color-aware-attr>] [remark-color-policy-attr>][remark-rcos-policy-attr>][<green-remark-rcos-attr>][<yellow-remark-rcos-attr>] [<base-etype-attr>] [<l4-src-port-attr>][<l4-dst-port-attr>][<src-ip-attr>][<dst-ip-attr>][<port-list-attr>][parent-profile <NAMESTRING(15)>]}.
```

where

<global-profile-create-object>	is the global profile l2-l3 being created.
<precedence-attr>	sets the precedence value for the profile from 0 to 255. The lower the number, the higher the precedence.
[dscp-attr-list>]	is the Differentiated Services Code Point (DSCP) value of matching frames.
[dot1dpri-attr-list>]	is the 802.1D priority value of matching frames.
[<vlan-id-attr-list>]	is the VLAN ID to monitor. This attribute supports up to four VLANs.
<cir-attr>	is the committed information rate in Kbps of this traffic meter, rounded to fit hardware.
<pir-attr>	is the peak information rate in Kbps of this traffic meter, rounded to fit hardware.
[<cbs-attr>]	is the meter committed burst size in Kbytes, rounded to nearest size supported in hardware.
[<ebs-attr>]	is the excess burst size in Kbps, rounded to nearest size supported in hardware.
[<dscp-remark-policy-attr>]	is the DSCP Remark Policy which can be leave or fixed. The fixed dscp value is configured by the fixed-dscp attribute.
[<fixed-dscp-attr>]	IPv4 frames that classify to the traffic profile have the DSCP value marked with a fixed value.
[<untagged-attr>]	is the untagged frame classifier.

Copyright© 2022 Ciena® Corporation

where

[<virtual-switch-attr>]	classifies the profile based on a virtual switch. Traffic profiles cannot classify to an MPLS VS.
[<ingress-color-aware-attr>]	indicates whether the profile is color aware or color blind.
[<remark-color-policy-attr>]	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 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© 2022 Ciena® Corporation

where

[<remark-rcos-policy-attr>]	is the meter output color-based RCoS remarking policy. Policies are: • leave: leaves the RCoS 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 RCoS/FCoS value to any packets that are classified as green. The new RCoS/FCos value is assigned using green-remark-rcos <NUMBER: 0-7>. • remark-yellow: assigns a new RCoS/FCoS value to any packets that are classified as yellow. The new RCoS/FCoS value is assigned using yellow-remark-rcos <NUMBER: 0-7>. • remark-both: assigns a new RCoS/FCoS value to any packets that are classified as green and to any packets that are classified as yellow. The new RCoS/FCoS value for green packets is assigned using green-remark-rcos <NUMBER: 0-7>. The new RCoS/FCoS value for yellow packets is assigned using yellow-remark-rcos <NUMBER: 0-7>. • remark-green-rcos-only: assigns a new RCoS value to any packets that are classified as green. The new RCoS value is assigned using green-remark-rcos <NUMBER:0-7>. • remark-yellow-rcos-only: assigns a new RCoS value to any packets that are classified as yellow. The new RCoS value is assigned using yellow-remark-rcos <NUMBER:0-7>. • remark-both-rcos-only: assigns a new RCoS value to any packets that are classified as green and to any packets that are classified as yellow. The new RCoS value for green packets is assigned using green-remark-rcos <NUMBER: 0-7>. The new RCoS value for yellow packets is assigned using yellow-remark-rcos <NUMBER: 0-7>. ***Note:*** This attribute cannot be set on hierarchical parent profiles.

Copyright© 2022 Ciena® Corporation

where

[<green-remark-rcos-attr>]	is the RCoS value for green remarking.
[<yellow-remark-rcos-attr>]	is the RCoS value for yellow remarking.
[<base-etype-attr>]	specifies the base Ethernet type to classify on. This classifier can be used to classify ARP packets.
[<l4-src-port-attr>]	specifies the l4 source port.
[<l4-dst-port-attr>]	specifies the l4 destination port.
[<src-ip-attr>]	specifies the source IPv4 address to classify on.
[<dst-ip-attr>]	specifies the destination IPv4 address to classify on.
[<port-list-attr>]	is the port list on the profile.
[parent-profile <NAMESTRING(15)>]	is the parent-profile name.

To modify attributes for a global traffic profile

2 Modify attributes for a global traffic profile:

```
traffic-profiling global-profile set <global-profile-
object-list>{[<precedence-attr>} [<dscp-attr-
list>][<dot1dpri-attr-list>][<vlan-id-attr-list>][<cir-
attr>][<pir-attr>][<cbs-attr>][<ebs-attr>][<dscp-remark-
policy-attr>][<fixed-dscp-attr>][<untagged-
attr>][<virtual-switch-attr>][<ingress-color-aware-
attr>][<remark-color-policy-attr>][<remark-rcos-policy-
attr>][<green-remark-rcos-attr>][<yellow-remark-rcos-
attr>][<base-etype-attr>][<l4-src-port-attr>][l4-dst-
port-attr>][<src-ip-attr>][<dst-ip-attr>][<port-list-
attr>][<ip-protocol-attr>]}
```

To remove classification attributes for a global traffic profile

3 Remove classification attributes for a global traffic profile:

```
traffic-profiling global-profile unset <global-profile-
object-list> {[dscp-attr-list][dot1dpri-attr-list][vlan-
attr-list][dscp-remark-policy][untagged][vs][ingress-
color-aware][remark-color-policy][remark-rcos-
policy>][base-etype][l4-src-port][l4-dst-port][src-
ip][dst-ip-attr][port][ip-protocol-attr][fixed-dscp-
value][yellow-remark-rcos][green-remark-rcos]}
```


Copyright© 2022 Ciena® Corporation

To delete a global traffic profile

4 Delete a global traffic profile:

```
traffic-profiling global-profile delete <global-profile-
object-list>
```

—end—

Example

This example creates a global profile and meter bases on BGP on vlan 10.

```
traffic-profiling global-profile create name BGP_V10 cir 1024
pir 2048 vlan 10 ip-protocol 6 l4-dst-port 179 precedence 4
```

This example also creates meter pcp value 3 on port 5, but not on BGP on vlan 10.

```
traffic-profiling global-profile create name pcp3 cir 1024 pir
2048 dot1dpri 3 port 5 precedence 5
```

This example creates a traffic profile named test1 that matches on packets of vlan 10 and dot1dpri 1, vlan 10 and dot1dpri2, vlan 20 and dot1dpri1, or vlan 20 and dot1dpri 2.

```
traffic-profiling global-profile create name test1 precedence
10 cir 1000 cbs 10 dot1dpri 1,2 vlan 10,20
```

This example creates a traffic profile named test2 that matches on packets of dscp 5, dscp 6, or dscp 7.

```
traffic-profiling global-profile create name test2 precedence
11 cir 1000 cbs 10 dscp 5-7
```


Copyright© 2022 Ciena® Corporation

Procedure 8-43
Configuring traffic profiling global profile parent

You can perform these operations on a global traffic profile parent:

- create
- modify
- remove classification attributes
- delete

Step	Action

To create a global traffic profile parent

1 Create a global traffic profile parent:

```
traffic-profiling global-profile parent create <global-
profile-create-object> {[<dscp-attr>][<dot1dpri-
attr>][<vlan-id-attr>][<cir-attr>][<cbs-attr>][<virtual-
switch-attr>][<base-etype-attr>][<l4-src-port-
attr>][<l4-dst-port-attr>][<src-ip-attr>][<dst-ip-
attr>][<port-list-attr>]}
```

where

<global-profile parent create-object>	is the global profile l2-l3 parent being created.
[dscp-attr>]	is the Differentiated Services Code Point (DSCP) value of matching frames.
[dot1dpri-attr>]	is the 802.1D priority value of matching frames.
[<vlan-id-attr>]	is the VLAN ID to monitor. This attribute supports up to one VLAN.
<cir-attr>	is the committed information rate in Kbps of this traffic meter, rounded to fit hardware.
[<cbs-attr>]	is the meter committed burst size in Kbytes, rounded to nearest size supported in hardware.
[<virtual-switch-attr>]	classifies the profile based on a virtual switch. Traffic profiles cannot classify to an MPLS VS.
<base-etype-attr>	specifies the base Ethernet type to classify on. This classifier can be used to classify ARP packets
<l4-src-port-attr>	specifies the l4 source port.
[<l4-dst-port-attr>]	specifies the l4 destination port.

Copyright© 2022 Ciena® Corporation

where

[<src-ip-attr>]	specifies the source IPv4 address to classify on.
[<dst-ip-attr>]	specifies the destination IPv4 address to classify on.
[<port-list-attr>]	is the port list on the profile.

To modify a global traffic profile parent

2 Create a global traffic profile parent:

```
traffic-profiling global-profile parent set <global-
profile-object> {[<cir-attr>][<cbs-attr>]}
```

To remove classification attributes for a global traffic profile parent

3 Remove classification attributes for a global traffic profile parent:

```
traffic-profiling global-profile parent unset <global-
profile-object> {[<cir-attr>][<cbs-attr>]}
```

To delete a global traffic profile parent

4 Delete a global traffic profile parent:

```
traffic-profiling global-profile parent delete <global-
profile-object>
```

—end—

Example

This example creates a parent profile that meters at 10000 Kbps.

```
traffic-profiling global-profile parent create name p1 cir
10000 cbs 10 vlan 10
```

This example shows parent profile, p1 with 2 children, ch1 and ch2. Profile ch1 matches packets coming in on port 5, with dscp 10 or dscp20. profile ch2 matches packets coming in on port 5, with vlan 100 or vlan 101.

```
traffic-profiling global-profile parent create name p1 cir
10000 cbs 10 port 5

traffic-profiling global-profile create name ch1 cir 1000 cbs
10 parent-profile p1 precedence 10 dscp 10,20

traffic-profiling global-profile create name ch2 cir 2000 cbs
10 parent-profile p1 precedence 11 vlan 100,101
```


Copyright© 2022 Ciena® Corporation

Procedure 8-44
Configuring per-port standard traffic profiling

This section shows how to configure per-port traffic profiling to meter ingress traffic based on 802.1D priority.

Step	Action
1	Enable traffic profiling globally. `traffic-profiling enable`
2	Enable traffic profiling: `traffic-profiling enable port <port>` where port <port> is the port number.
3	Confirm the global configuration (optional). `traffic-profiling show`
4	Set traffic profiling mode on the specified port to standard-dot1dpri (default) and a narrow classifier-mode. `traffic-profiling set <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>} {classifier-mode <narrow|wide}`

Copyright© 2022 Ciena® Corporation

where

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.
classifier mode <narrow | wide>	is the classifier mode. Narrow excludes dropped frames from meters. Wide does not.

5 Confirm the per port configuration:

```
traffic-profiling show port <port>
```

6 Create a traffic profile on a port and specify CIR, PIR and dot1dpri values:

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>} cir <NUMBER>
pir <NUMBER> dot1dpri <NUMBER>
```

where

profile <NUMBER: 1..4096> | name <STRING>	is the traffic profile number or name.
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.
dot1dpri <NUMBER>	is the .1D priority of matching frames.

7 Create a second traffic profile on the same port and specify CIR, PIR and dot1dpri:

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>}{cir <NUMBER>}
{pir <NUMBER>} {dot1dpri <NUMBER>}
```


Copyright© 2022 Ciena® Corporation

8 Create a third traffic profile on the same port and specify CIR, PIR and dot1dpri:

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>} {cir
<NUMBER>} {pir <NUMBER>} {dot1dpri <NUMBER>}
```

9 Create a fourth traffic profile on the same port and specify CIR and PIR for untagged traffic.

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>} {cir
<NUMBER>} {pir <NUMBER>} [untagged]
```

where

untagged the profile matches on untagged frames.

10 Display the traffic profiles associated with port 1 (Optional).

```
traffic-profiling standard-profile show <port>
```

—end—

Example

This example shows how to configure per-port traffic profiling to meter ingress traffic based on 802.1D priority.

```
traffic-profiling enable
traffic-profiling enable port 1
traffic-profiling show
traffic-profiling set port 1 mode standard-dot1dpri classifier-mode narrow
traffic-profiling show port 1
traffic-profiling standard-profile create port 1 profile 1 cir 3000 pir 3200
dot1dpri 0,1,2
traffic-profiling standard-profile create port 1 profile 2 cir 3000 pir 3200
dot1dpri 3,4,5
traffic-profiling standard-profile create port 1 profile 3 cir 2000 pir 2432
dot1dpri 6,7
traffic-profiling standard-profile create port 1 profile 4 cir 128 pir 256
untagged
traffic-profiling standard-profile show port 1
```


Copyright© 2022 Ciena® Corporation

Procedure 8-45
Configuring per-port and per-VLAN standard traffic profiling

Configure per-port and per-VLAN traffic profiling to meter ingress traffic based on VLAN tag.

Step	Action
1	Enable traffic profiling globally. `traffic-profiling enable`
2	Enable traffic profiling: `traffic-profiling enable port <port>` where port <port> is the port number.
3	Confirm the global configuration (optional). `traffic-profiling show`
4	Set traffic profiling mode on the specified port to standard-dot1dpri (default) and a narrow classifier-mode. `traffic-profiling set <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>} {classifier-mode <narrow|wide}` where 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. classifier mode <narrow | wide> is the classifier mode. Narrow excludes dropped frames from meters. Wide does not.

Copyright© 2022 Ciena® Corporation

5 Confirm the per port configuration for the specified port.

```
traffic-profiling show port <port>
```

6 Create the VLAN:

```
vlan create vlan <vlan>
```

where

vlan <vlan>	is the VLAN(s) being created.

7 Add the specified ports to the VLANs created:

```
vlan add vlan <vlan> port <port>
```

8 Confirm the VLAN configuration (optional).

```
vlan show
```

9 Create a traffic profile on a port and specify CIR, PIR and VLAN values:

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>} {cir
<NUMBER>} {pir <NUMBER>} {classifier vlan <vlan>}
```

where

profile <NUMBER: 1..4096> I name <STRING>	is the traffic profile number or name.
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.
classifier vlan <vlan>	is the VLAN that was created.

10 Create a second traffic profile on the same port and specify CIR, PIR and classifier vlan values:

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>} cir <NUMBER>
pir <NUMBER> classifier vlan <vlan>
```

11 Create a third traffic profile on the same port and specify CIR, PIR and classifier VLAN:

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>} cir <NUMBER>
pir <NUMBER> classifier vlan <vlan>
```

12 Create a fourth traffic profile on the same port and specify CIR and PIR for untagged traffic:

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>} untagged cir
<NUMBER> pir <NUMBER>
```


Copyright© 2022 Ciena® Corporation

13 Display the traffic profiles associated with port 1 (Optional).

```
traffic-profiling standard-profile show port <port>
```

—end—

Example

This example shows how to configure per-port and per-VLAN standard traffic profiling.

```
traffic-profiling enable
traffic-profiling enable port 1
traffic-profiling show
traffic-profiling set port 1 mode standard-vlan classifier-mode narrow
traffic-profiling show port 1
vlan create vlan 1001-1003

...This may take a while depending upon the VLAN range...
vlan add vlan 1001-1003 port 1, 9

...This may take a while depending upon the VLAN range...
vlan show
traffic-profiling standard-profile create port 1 profile 1 cir 100 pir 1024
vlan 1001
traffic-profiling standard-profile create port 1 profile 2 cir 100 pir 1024
vlan 1002
traffic-profilng standard-profile create port 1 profile 3 cir 100 pir 1024
vlan 1003
traffic-profiling standard-profile create port 1 profile 4 untagged cir 128
pir 256
traffic-profiling standard-profile show port 1
```


Copyright© 2022 Ciena® Corporation

Procedure 8-46
Configuring hierarchical VLAN standard traffic profiles

Configure hierarchical VLAN standard traffic profiles.

Step	Action
1	Set the port mode to hierarchical VLAN mode. `traffic-profiling set port <port> mode hierarchical-vlan` where port <port> — is the port number.
2	Create the parent profile to classify traffic with a VLAN with a narrow classifier-mode and specify the child mode to classify based on 802.1D priority: `traffic-profiling standard-profile create port <port> {profile <NUMBER: 1..4096> | name <STRING>} child-mode standard-dot1dpri {cir <NUMBER>} {cbs <#>} {vlan <vlan>} classifier-mode narrow` where profile <NUMBER: 1..4096> | name <STRING> — is the parent traffic profile number or name. cir <NUMBER> — is the committed 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. vlan <vlan> — is the VLAN-ID to monitor.

Copyright© 2022 Ciena® Corporation

3 Create the child profile to classify traffic with 802.1D priority values:

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>} {parent
<Traffic Profiling Standard Name>}{cir <NUMBER>}{pir
<NUMBER>}{cbs <#>}{ebs <#>}
```

where

profile <NUMBER: 1..4096> | name <STRING>	is the traffic profile number or name.
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 meter excess burst size in Kbytes, rounded to nearest size supported in hardware.

4 Create a child default under the VLAN:

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>} {cir <Kbps>}
{pir <NUMBER>} {cbs <#>} {ebs <#> name <vlan>} {parent
<Traffic Profiling Standard Name>}
```

where

profile <NUMBER: 1..4096> | name <STRING>	is the default profile number or name.

5 Create the parent profile to classify traffic with the VLAN with a narrow classifier-mode and specify the child mode to classify based on DSCP priority.

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>} child-mode
standard-dscp {cir <NUMBER>} {cbs <#>} {vlan <vlan>}
```


Copyright© 2022 Ciena® Corporation

6 Create the child profile to classify traffic with a specified DSCP value:

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>} {parent
<Traffic Profiling Standard Name>} {dscp <dscp values and
ranges>} {cir <NUMBER>} {pir <NUMBER>} {cbs <#>}{ebs <#>}
```

where

dscp <dscp values and ranges>	is the DSCP value of the matching frames.

7 Create the child profile to classify traffic with a specified DSCP value:

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>} {parent
<Traffic Profiling Standard Name>} {dscp <dscp values and
ranges>} {cir <NUMBER>} {pir <NUMBER>} {cbs <#>}{ebs <#>}
```

8 Create the child profile to classify and drop untagged traffic.

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>} {parent
<Traffic Profiling Standard Name>} drop true untagged
```

9 Create the parent profile to classify traffic with the VLAN with a narrow classifier-mode and specify the child mode to classify based on IP precedence;

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>} child-mode
standard-ip-prec {cir <NUMBER>} {cbs <#>} {vlan <VLAN>}
```

10 Create the child profile to classify traffic with an IP precedence value:

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>} {parent
<Traffic Profiling Standard Name>} child-mode ip-prec
<NUMBER> {cir <NUMBER>}{pir <NUMBER>{{cbs <#>}{ebs <#>}
```

11 Create the child profile to classify and drop untagged traffic.

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1..4096> | name <STRING>}-untagged
{parent <Traffic Profiling Standard Name>} drop true
untagged
```

Note: Alternately, an profile can be created at the same level as the parent to classify untagged traffic.

12 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 <port>
{profile <NUMBER: 1..4096> | name <STRING>} default-
profile true {cir <NUMBER>} {pir <NUMBER>} {cbs <#>} {ebs
<#>} {name <String>}
```

—end—

Copyright© 2022 Ciena® Corporation

Example

This example shows how to configure hierarchical VLAN traffic profiles.

```
traffic-profiling set port 6 mode hierarchical-vlan
traffic-profiling standard-profile create port 6 name parent1 child-mode
standard-dot1dpri cir 50000 cbs 8 vlan 2000 classifier-mode narrow
traffic-profiling standard-profile create port 6 name child1 parent parent1
dot1dpri 0,1 cir 30000 pir 50000 cbs 4 ebs 2
traffic-profiling standard-profile create port 6 name child2 parent parent1
dot1dpri 2 cir 10000 pir 50000 cbs 4 ebs 2
traffic-profiling standard-profile create port6 default-profile true cir 64
pir 45000 cbs 4 ebs 4 name vlan2000child parent parent1
```

Figure 8-22
Hierarchical VLAN with 802.1D priority

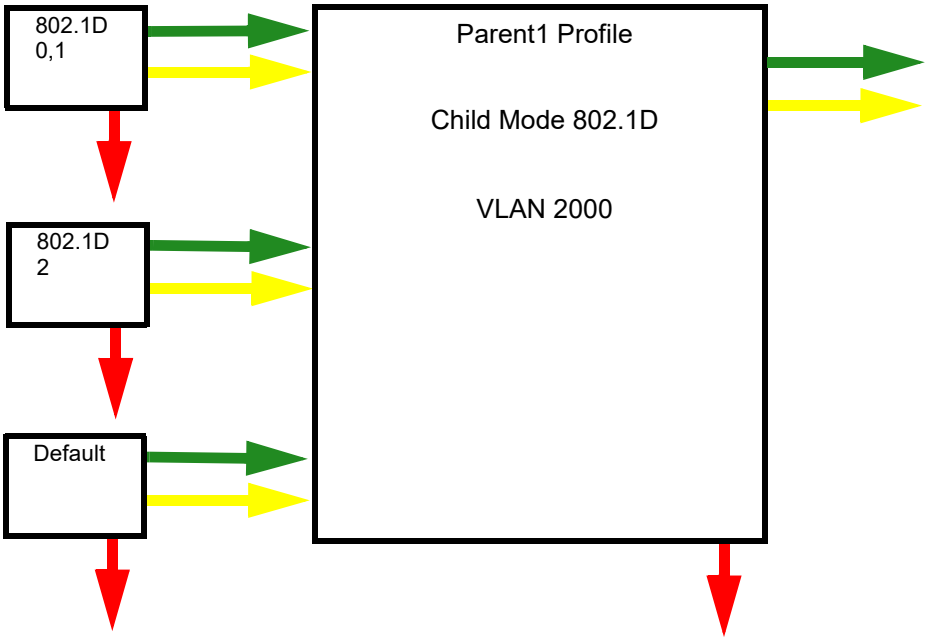


```
traffic-profiling standard-profile create port 6 name parent2 child-mode
standard-dscp cir 50000 cbs 8 vlan 4000
traffic-profiling standard-profile create port 6 name child1dscp parent
parent2 dscp 8 cir 20000 pir 50000 cbs 4 ebs 32
traffic-profiling standard-profile create port 6 name child2dscp parent
parent2 dscp 16 cir 15000 pir 50000 cbs 4 ebs 16
traffic-profilng standard-profile create port 6 name child3dscp parent parent2
dscp 32 cir 14000 pir 50000 cbs 4 ebs 16
traffic-profiling standard-profile create port 1 name child4untagged parent
parent2 drop true untagged
```


Copyright© 2022 Ciena® Corporation

Figure 8-23
Hierarchical VLAN with DSCP priority

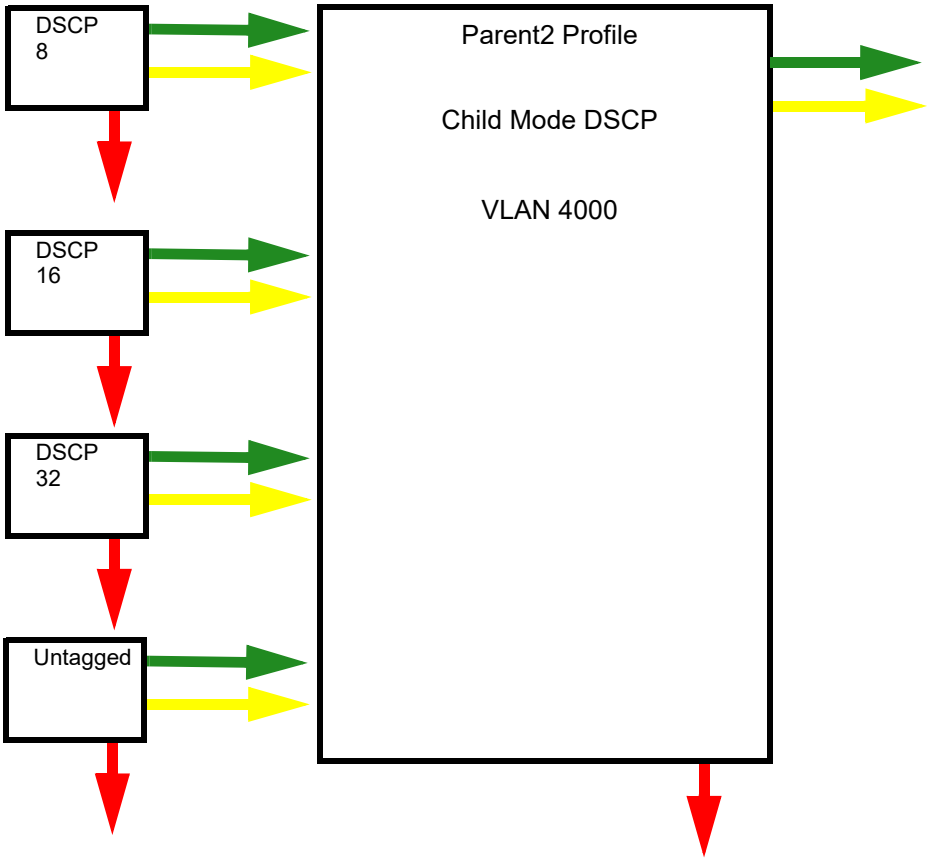


```
traffic-profiling standard-profile create port 6 name parentIpp child-mode
standard-ip-prec cir 50000 cbs 8 vlan 2100
traffic-profiling standard-profile create port 6 name child1ipp parent
parentIpp ip-prec 6 cir 20000 pir 50000 cbs 4 ebs 16
traffic-profiling standard-profile create port 6 name child2ipp parent
parentIpp ip-prec 5 cir 10000 pir 50000 cbs 4 ebs 16
traffic-profiling standard-profile create port 6 name child3ipp parent
parentIPP ip-prec 4 cir 19000 pir 50000 cbs 4 ebs 16
traffic-profiling standard-profile create port 1 name child4ipp-untagged
parent parentIpp drop true untagged
```


Copyright© 2022 Ciena® Corporation

Figure 8-24
Hierarchical VLAN with IP precedence

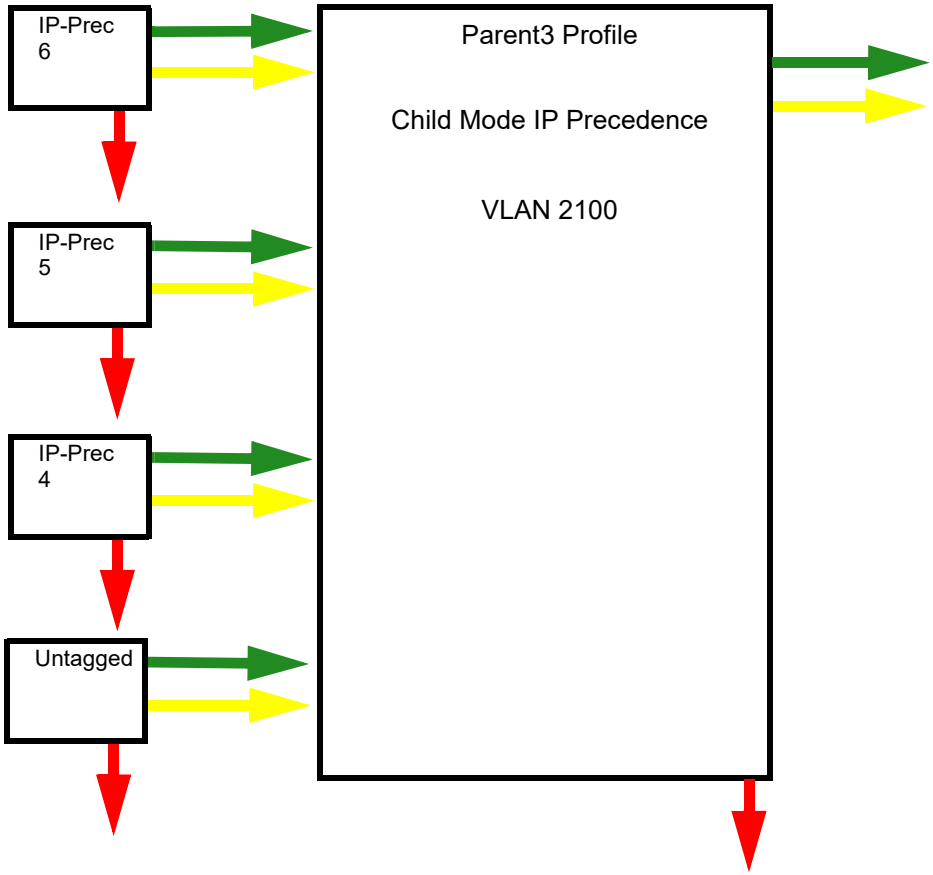


```
traffic-profiling standard-profile create port 6 default-profile true cir
15000 pir 45000 cbs 8 ebs 8 name Tidefault
```


Copyright© 2022 Ciena® Corporation

Procedure 8-47
Configuring hierarchical standard port traffic profiles

Configure hierarchical standard port traffic profiles.

Step	Action
1	Set the port mode to hierarchical port mode. `traffic-profiling set port <port> 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 <port> {profile <NUMBER: 1..4096> | name <STRING>} child-mode standard-dot1dpri {cir <NUMBER>} {cbs <#>}` where port <port> — is the name of the port. profile <NUMBER: 1..4096> | name <STRING> — is the traffic profile number or name. cir <NUMBER> — is the committed 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.
3	Create the child profile to classify traffic with 802.1D priority values. `traffic-profiling standard-profile create port <port> {profile <NUMBER: 1..4096> | name <STRING>} {parent <Traffic Profiling Standard Name>} {dot1dpri <NUMBER>} {cir <NUMBER>}{pir <NUMBER>} {cbs <#>} {ebs <#>}` where parent <Traffic Profiling Standard Name> — is the parent traffic profile number. dot1dpri <NUMBER> — is the .1D priority of matching frames. ebs <Kbytes> — is the meter excess burst size in KBytes, rounded to nearest size supported in hardware.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example shows how to configure one parent profile with the 802.1D child mode and three child profiles for classification of specific 802.1D values.

```
traffic-profiling set port 5 mode hierarchical-port
traffic-profiling standard-profile create port 5 name portdot1D child-mode
standard-dot1dpri cir 50000 cbs 8
traffic-profiling standard0profile create port 5 name childdot1D1 parent
portdot1D dot1dpri 6 cir 30000 pir 4000 cbs 4 ebs 2
traffic-profiling standard-profile create port 5 name childdot1D2 parent port
dot1D dot1dpri 3 cir 10000 pir 2500 cbs 4 ebs 8
traffic-profiling standard-profile create port 5 name childdot1D3 parent
portdot1D dot1dpri 4 cir 9800 pir 9800 cbs 4 ebs 16
```

This figure describes the above example.

Figure 8-25
Hierarchical port with 802.1D priority

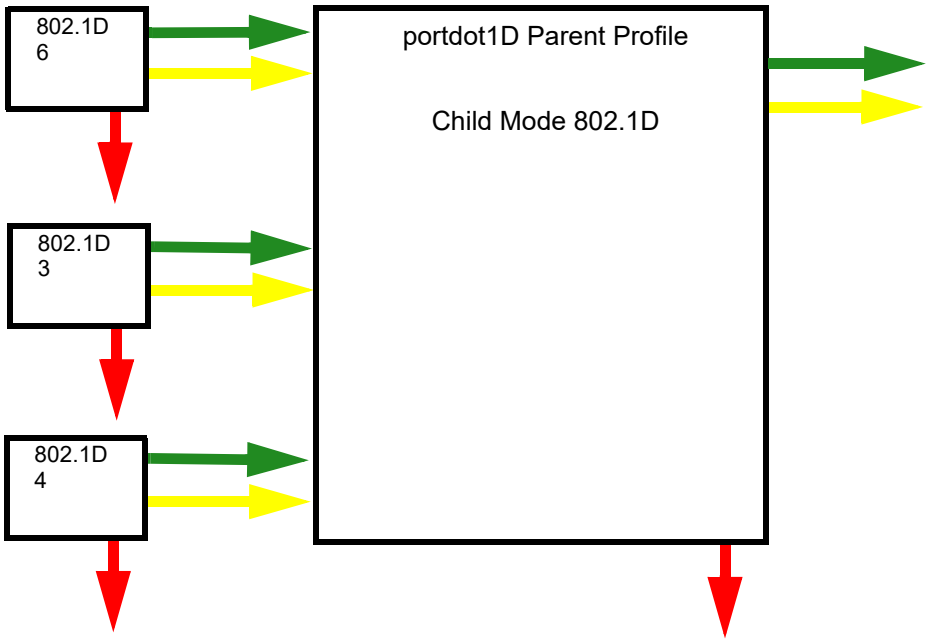

Copyright© 2022 Ciena® Corporation

Procedure 8-48
Configuring VS classification for standard traffic profiles

Configure VS classification for standard traffic profiles.

Step	Action

1	Create the virtual switches and add them to the respective VLANs.
	`virtual-switch add reserved-vlan <VLANID List>`
	`virtual-switch ethernet create vs <vs>`
	`virtual-switch ethernet add vs <vs> port <port list> vlan <VLAN List>`
	where
	reserved-vlan <VLAN ID List> — is the VLAN list.
	vs <vs> — is the virtual switch name.
	port <port list> — is the subscriber port list

2	Enable traffic profiling and set the mode to advanced.
	`traffic-profiling enable`
	`traffic-profiling enable port <port>`
	`traffic-profiling set port <port> mode advanced`
	where
	port <port> — is the port number.

3	Create a traffic profile and assign it to a VS, specifying CIR and PIR values. `traffic-profiling standard-profile create port <port> {profile <NUMBER: 1-4096> | name <STRING>} {cir <NUMBER>} {pir <NUMBER> {vs <Virtual Switch Ethernet Name>}` where profile <NUMBER: 1-4096> | name <STRING> — is the traffic profile number or name.

Copyright© 2022 Ciena® Corporation

where

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.
vs <Virtual Switch Ethernet Name>	is the virtual switch as a classifier (EVC matching frames).

4 Create traffic profiles and assign them to the VS, specifying CIR, PIR, VLAN, and 802.1D values.

```
traffic-profiling standard-profile create port <port>
{profile <NUMBER: 1-4096> | name <STRING>} {cir <NUMBER>}
{pir <NUMBER>} {vs <Virtual Switch Ethernet Name>} {vlan
<VLAN>}{dot1dpri <NUMBER>}
```

where

dot1dpri <NUMBER>	is the .1D priority of matching frames.

—end—

Example

In this configuration example, 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.

```
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
virtual0switch ethernet add vs VS3 port 1 vlan 3002
traffic-profiling enable
traffic-profiling enable port 1
traffic-profiling set port 1 mode advanced
traffic-profiling standard-profile create port 1 profile 1 cir 1024 pir 2048
vs VS1
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 4 cir 2048 pir 2048
vs VS2 vlan 2002 dot1dpri 6,7
```


Copyright© 2022 Ciena® Corporation

Procedure 8-49
Setting the VLAN range for standard traffic profiles

Set the vlan-range for standard profiles as required by the network plan.

The following notes apply:

- There is one individual traffic profile for each EVPL.
- It is recommended to use traffic profiles for those EVPLs that are added with single CVID values. Otherwise, their traffic will be metered by the same traffic profile that is configured for the EVPL with a vlan-range and having encap-fixed-dot1dpri the same as the resolved cos derived for single CVID EVPL.
- All the EVPLs with a vlan-range must have a unique encap-fixed-dot1dpri value and corresponding traffic profile also.
- Traffic profiles can be used to remark PCP bits of frames, if needed.

Step	Action
1	Create a VLAN: `vlan create vlan <VlanIdList>`
2	Add ports to the VLAN created. `vlan add vlan <VlanList> port <PortNameList>`
3	Create a static Ethernet virtual circuit for the specified VLAN object: `virtual-circuit ethernet create vc <vc_name> vlan <VlanID>` where vc <vc> is the virtual circuit name. vlan <vlan name> is the VLAN name or vlan ID.
4	Create a virtual switch instance and associate a virtual circuit: `virtual-switch ethernet create vs <vs> vc <vc>`

Copyright© 2022 Ciena® Corporation

5 Add UNI memberships to the virtual switch.

```
virtual-switch ethernet add vs <vs> port <port list> vlan
<VLAN List>
```

where

vlan <VLAN ID List>	is the VLAN list.
vs <vs>	is the virtual switch name.
port <port list>	is the subscriber port list

6 Configure the cos policy for the vlan.

```
encap-cos-policy <dot1dpri-inherit | fixed>
```

where

dot1dpri <NUMBER>	is the .1D priority of matching frames.

7 Configure the vlan-range type:

```
vlan-range-type <strict|group>
```

where

vlan-range-type <strict | group>	is the vlan-range type.

8 Configure the encap cos policy for each vlan-range.

```
encap-cos-policy <dot1dpri-inherit | fixed>
```

where

dot1dpri <NUMBER>	is the 1D priority of matching frames.

9 Configure a standard traffic profile for EVPLs with vlan-range as group:

```
traffic-profiling standard-profile create port <port> vs
<virtual switch name> cir <number> cbs <number> profile
<profile number or name> vlan <vlan id or vlan range>
```

—end—

Copyright© 2022 Ciena® Corporation

Example

In this configuration example, there are traffic profiles for EVPLs with a single CVID. All EVPLs with a range have a unique encap-dot1d-pro value and corresponding traffic profile.

```
vlan create vlan 100
vlan add vlan 100 port 2
virtual-circuit eth create vc vc1 vlan 100
virtual-switch eth create vs vs1 vc vc1
virtual-switch ethernet add vs vs1 port 1 vlan 10 encap-cos-policy fixed
encap-fixed-dot1dpri 0
virtual-switch ethernet add vs vs1 port 1 vlan 20 encap-cos-policy fixed
encap-fixed-dot1dpri 1
virtual-switch ethernet add vs vs1 port 1 vlan 100-199 vlan-range-type strict
encap-cos-policy fixed encap-fixed-dot1dpri 0
virtual-switch ethernet add vs vs1 port 1 vlan 200-299 vlan-range-type strict
encap-cos-policy fixed encap-fixed-dot1dpri 1
virtual-switch ethernet add vs vs1 port 1 vlan 300-399 vlan-range-type strict
encap-cos-policy fixed encap-fixed-dot1dpri 2
virtual-switch ethernet add vs vs1 port 1 vlan 400-499 vlan-range-type strict
encap-cos-policy fixed encap-fixed-dot1dpri 3
virtual-switch ethernet add vs vs1 port 1 vlan 500-599 vlan-range-type strict
encap-cos-policy fixed encap-fixed-dot1dpri 4
virtual-switch ethernet add vs vs1 port 1 vlan 600-699 vlan-range-type strict
encap-cos-policy fixed encap-fixed-dot1dpri 5
virtual-switch ethernet add vs vs1 port 1 vlan 700-799 vlan-range-type strict
encap-cos-policy fixed encap-fixed-dot1dpri 6
virtual-switch ethernet add vs vs1 port 1 vlan 800-899 vlan-range-type strict
encap-cos-policy fixed encap-fixed-dot1dpri 7

traffic-profiling set port 1 mode advanced
traffic-profiling enable
traffic-profiling standard-profile create port 1 vs vs1 cir 1024 cbs 256
profile 1 name p1 vlan 10
traffic-profiling standard-profile create port 1 vs vs1 cir 1024 cbs 256
profile 2 name p2 vlan 20
traffic-profiling standard-profile create port 1 vs vs1 cir 1024 cbs 256
profile 3 name p3 vlan 100-199
traffic-profiling standard-profile create port 1 vs vs1 cir 1024 cbs 256
profile 4 name p4 vlan 200-299
traffic-profiling standard-profile create port 1 vs vs1 cir 1024 cbs 256
profile 5 name p5 vlan 300-399
traffic-profiling standard-profile create port 1 vs vs1 cir 1024 cbs 256
profile 6 name p6 vlan 400-499
traffic-profiling standard-profile create port 1 vs vs1 cir 1024 cbs 256
profile 7 name p7 vlan 500-599
traffic-profiling standard-profile create port 1 vs vs1 cir 1024 cbs 256
profile 8 name p8 vlan 600-699
traffic-profiling standard-profile create port 1 vs vs1 cir 1024 cbs 256
profile 9 name p9 vlan 700-799
traffic-profiling standard-profile create port 1 vs vs1 cir 1024 cbs 256
profile 10 name p10 vlan 800-899
```


Copyright© 2022 Ciena® Corporation

Procedure 8-50
Configuring standard traffic profiling for EVPLs with the vlan-range-type as group

Configure standard traffic profiling for EVPLs with vlan-range-type as group.

Step	Action
1	Configure a standard traffic profile for EVPLs with vlan-range as group: `traffic-profiling standard-profile create port <port> vs <virtual switch name> cir <number> cbs <number> profile <profile number or name> vlan <vlan id or vlan range>`

—end—

Example

In this configuration example, a single UNI port has UNI memberships of two Q-in-Q VS's. All the configuration options to optimize available resources have been used. Port + VS + VLAN classifiers of traffic profiling are used to ensure per range per VS traffic profile configuration.

```
vlan create vlan 1000
vlan add vlan 1000 port 2
virtual-circuit eth create vc vc2 vlan 1000 vlan-transform i-push,e-pop
virtual-switch eth create vs vs2 vc vc2
virtual-switch ethernet add vs vs2 port 1 vlan 3000 encap-cos-policy fixed
encap-fixed-dot1dpri 0
virtual-switch ethernet add vs vs2 port 1 vlan 3001 encap-cos-policy fixed
encap-fixed-dot1dpri 1
port set port 2 vlan-range-index 1 vlan-range 1000-1099
virtual-switch ethernet add vs vs2 port 1 vlan 1000-1099 vlan-range-type group
encap-cos-policy fixed encap-fixed-dot1dpri 0
port set port 2 vlan-range-index 2 vlan-range 1100-1199
virtual-switch ethernet add vs vs2 port 1 vlan 1100-1199 vlan-range-type group
encap-cos-policy fixed encap-fixed-dot1dpri 1
port set port 2 vlan-range-index 3 vlanr-range 1200-1299
virtual-switch ethernet add vs vs2 port 1 vlan 1200-1299 vlan-range-type group
encap-cos-policy fixed encap-fixed-dot1dpri 2
port set port 2 vlan-range-index 4 vlan-range 1300-1399
virtual-switch ethernet add vs vs2 port 1 vlan 1300-1399 vlan-range-type group
encap-cos-policy fixed encap-fixed-dot1dpri 3
port set port 2 vlan-range-index 5 vlan-range 1400-1499
virtual-switch ethernet add vs vs2 port 1 vlan 1400-1499 vlan-range-type group
encap-cos-policy fixed encap-fixed-dot1dpri 4
port set port 2 vlan-range-index 6 vlan-range 1500-1599
virtual-switch ethernet add vs vs2 port 1 vlan 1500-1599 vlan-range-type group
encap-cos-policy fixed encap-fixed-dot1dpri 5
port set port 2 vlan-range-index 7 vlan-range 1600-1699
virtual-switch ethernet add vs vs2 port 1 vlan 1600-1699 vlan-range-type group
encap-cos-policy fixed encap-fixed-dot1dpri 6
port set port 2 vlan-range-index 8 vlan-range 1700-1799
```


Copyright© 2022 Ciena® Corporation

```
virtual-switch ethernet add vs vs2 port 1 vlan 1700-1799 vlan-range-type group
encap-cos-policy fixed encap-fixed-dot1dpri 7

traffic-profiling standard-profile create port 1 vs vs2 cir 1024 cbs 256
profile 11 name p11 vlan 3000
traffic-profiling standard-profile create port 1 vs vs2 cir 1024 cbs 256
profile 12 name p12 vlan 3001
traffic-profiling standard-profile create port 1 vs vs2 cir 1024 cbs 256
profile 13 name p13 vlan 1000-1099
traffic-profiling standard-profile create port 1 vs vs2 cir 1024 cbs 256
profile 14 name p14 vlan 1100-1199
traffic-profiling standard-profile create port 1 vs vs2 cir 1024 cbs 256
profile 15 name p15 vlan 1200-1299
traffic-profiling standard-profile create port 1 vs vs2 cir 1024 cbs 256
profile 16 name p16 vlan 1300-1399
traffic-profiling standard-profile create port 1 vs vs2 cir 1024 cbs 256
profile 17 name p17 vlan 1400-1499
traffic-profiling standard-profile create port 1 vs vs2 cir 1024 cbs 256
profile 18 name p18 vlan 1500-1599
traffic-profiling standard-profile create port 1 vs vs2 cir 1024 cbs 256
profile 19 name p19 vlan 1600-1699
traffic-profiling standard-profile create port 1 vs vs2 cir 1024 cbs 256
profile 20 name p20 vlan 1700-1799
```


Copyright© 2022 Ciena® Corporation

Procedure 8-51
Configuring a global profile with the VLAN range option

Configure global traffic profiles for EVPLs with the vlan-range-type as strict.

- Only a single UNI port can be added as a classifier in global profiles having VLAN list as classifier.
- Global profiles for EVPLs with a single CVID should have lower precedence than global profiles configured for EVPLs with a VLAN range.

Step	Action

1	Configure a global profile with the vlan-range-type as strict:
	`traffic-profiling global-profile create <global-profilecreate-object> {<precedence-attr> [<dscp-attr-list>][<dot1dpri-attr-list>] [<vlan-id-attr-list>] <cir-attr> <pir-attr> [<cbs-attr>] [<ebs-attr>] [<dscp-remarkpolicy-attr>] [<fixed-dscp-attr>][<untagged-attr>]`
	`[<virtual-switch-attr>] [<ingress-color-aware-attr>][remark-color-policy-attr>][remark-rcos-policyattr>][<green-remark-rcos-attr>][<yellow-remark-rcosattr>][<base-etype-attr>] [<l4-src-port-attr>][<l4-dstport-attr>][<src-ip-attr>][<dst-ip-attr>][<port-listattr>][parent-profile <NAMESTRING(15)>]}`

Copyright© 2022 Ciena® Corporation

Example

In this configuration example, Port 1 has multiple UNI memberships. There are single CVID EVPLs and EVPLs with vlan-range-type configurations. Global profiling configuration is provided for all the EVPLs.

```
vlan create vlan 100
vlan add vlan 100 port 2
virtual-circuit ethernet create vc vc1 vlan 100
virtual-switch ethernet create vs vs1 vc vc1
virtual-switch ethernet add vs vs1 port 1 vlan 10 encap-cos-policy fixed
encap-fixed-dot1dpri 0
virtual-switch ethernet add vs vs1 port 1 vlan 20 encap-cos-policy fixed
encap-fixed-dot1dpri 1
virtual-switch ethernet add vs vs1 port 1 vlan 100-199 vlan-range-type strict
encap-cos-policy fixed encap-fixed-dot1dpri 0
virtual-switch ethernet add vs vs1 port 1 vlan 200-299 vlan-range-type strict
encap-cos-policy fixed encap-fixed-dot1dpri 1
virtual-switch ethernet add vs vs1 port 1 vlan 300-399 vlan-range-type strict
encap-cos-policy fixed
virtual-switch ethernet add vs vs1 port 1 vlan 400-499 vlan-range-type strict
encap-cos-policy fixed encap-fixed-dot1dpri 3
virtual-switch ethernet add vs vs1 port 1 vlan 500-599 vlan-range-type strict
encap-cos-policy fixed encap-fixed-dot1dpri 4
virtual-switch ethernet add vs vs1 port 1 vlan 600-699 vlan-range-type strict
encap-cos-policy fixed encap-fixed-dot1dpri 5
virtual-switch ethernet add vs vs1 port 1 vlan 700-799 vlan-range-type strict
encap-cos-policy fixed encap-fixed-dot1dpri 6
virtual-switch ethernet add vs vs1 port 1 vlan 800-899 vlan-range-type strict
encap-cos-policy fixed encap-fixed-dot1dpri 7
traffic-profiling set profile-mode global-l2-l3
traffic-profiling global-profile create name p1 cir 1024 pir 1024 cbs 256 ebs
0 precedence 1
traffic-profiling global-profile set profile p1 vlan 10
traffic-profiling global-profile set profile p1 port 1
traffic-profiling global-profile set profile p1 remark-rcos-policy remark-
both green-remark-rcos 5 yellow-remark-rcos 5
traffic-profiling global-profile enable profile p1
traffic-profiling global-profile create name p2 cir 1024 pir 1024 cbs 256 ebs
0 precedence 2
traffic-profiling global-profile set profile p2 vlan 20
traffic-profiling global-profile set profile p2 port 1
traffic-profiling global-profile enable profile p2
traffic-profiling global-profile create name p3 cir 1024 pir 1024 cbs 256 ebs
0 precedence 3
traffic-profiling global-profile set profile p3 vlan 100-199
traffic-profiling global-profile set profile p3 port 1
traffic-profiling global-profile enable profile p3
traffic-profiling global-profile create name p4 cir 1024 pir 1024 cbs 256 ebs
0 precedence 4
traffic-profiling global-profile set profile p4 vlan 200-299
traffic-profiling global-profile set profile p4 port 1
traffic-profiling global-profile enable profile p4
traffic-profiling global-profile create name p5 cir 1024 pir 1024 cbs 256 ebs
0 precedence 5
traffic-profiling global-profile set profile p5 vlan 300-399
traffic-profiling global-profile set profile p5 port 1
traffic-profiling global-profile enable profile p5
```


Copyright© 2022 Ciena® Corporation

```
traffic-profiling global-profile create name p6 cir 1024 pir 1024 cbs 256 ebs
0 precedence 6
traffic-profiling global-profile set profile p6 vlan 400-499
traffic-profiling global-profile set profile p6 port 1
traffic-profiling global-profile enable profile p6
traffic-profiling global-profile create name p7 cir 1024 pir 1024 cbs 256 ebs
0 precedence 7
traffic-profiling global-profile set profile p7 vlan 500-599
traffic-profiling global-profile set profile p7 port 1
traffic-profiling global-profile enable profile p7
traffic-profiling global-profile create name p8 cir 1024 pir 1024 cbs 256 ebs
0 precedence 8
traffic-profiling global-profile set profile p8 vlan 600-699
traffic-profiling global-profile set profile p8 port 1
traffic-profiling global-profile enable profile p8
traffic-profiling global-profile create name p9 cir 1024 pir 1024 cbs 256 ebs
0 precedence 9
traffic-profiling global-profile set profile p9 vlan 700-799
traffic-profiling global-profile set profile p9 port 1
traffic-profiling global-profile enable profile p9
traffic-profiling global-profile create name p10 cir 1024 pir 1024 cbs 256 ebs
0 precedence 10
traffic-profiling global-profile set profile p10 vlan 800-899
traffic-profiling global-profile set profile p10 port 1
traffic-profiling global-profile enable profile p10
traffic-profiling enable
```


Copyright© 2022 Ciena® Corporation

Procedure 8-52
Configuring VS classification for standard HIM traffic profiles

Configure VS classification for standard HIM traffic profiles.

Step	Action
1	Create the virtual switch configuration as described in Procedure 8-48.

2	Enable traffic profiling and set the mode to advanced.
	`traffic-profiling enable`
	`traffic-profiling enable port <port>`
	`traffic-profiling set port <port> mode advanced`
	where
	port <port> is the port number.

3	Create a parent traffic profile and assign it to the VS with child-mode VLAN-CoS, specifying CIR and PIR values. `traffic-profiling standard-profile create port <port> {profile <NUMBER: 1..4096> | name <STRING>} child-mode vlan-cos {cir <NUMBER>} {vs <Virtual Switch Ethernet Name>}` where profile <NUMBER: 1..4096> | name <STRING> is the traffic profile number or name cir <NUMBER> is the committed information rate in Kbps of this traffic meter, rounded to fit hardware. vs <Virtual Switch Ethernet Name> is the Virtual Switch as a classifier (EVC matching frames)

Copyright© 2022 Ciena® Corporation

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 <port>
{profile <NUMBER: 1..4096> | name <STRING>} {parent
<Traffic Profiling Standard Name>} {cir <NUMBER>} {pir
<NUMBER>} {dscp <dscp values and ranges>}
```

where

parent <Traffic Profiling Standard Name>	is the parent traffic profile number.
dscp <dscp values and ranges>	is the DSCP value of the matching frames.

—end—

Example

In this configuration example, regardless of the VLAN, any traffic ingressing VS3 with the specified CoS values is classified/metered.

```
traffic-profiling enable
traffic-profiliing enable port 1
traffic-profiling set port 1 mode advanced
traffic-profiling standard-profile create port 1 name Parent1 child-mode vlan-
cos cir 50048 vs VS3
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
```


Copyright© 2022 Ciena® Corporation

Procedure 8-53
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

port <PortName> is the port that you want to display associated standard traffic profiles for.

—end—

Example

This example shows sample output for displaying all 392X standard traffic profiles.

```
> traffic-profiling standard-profile show
+--------------------------- STANDARD PROFILE TABLE --------------------------+
|Port|ID|Profile    |     |Parent|  BW (Kbps)  |Max Burst KB|CLASSIFIERS|Resource|
|    |  |Name       |Role|#Child|CIR   |PIR   |CBS|    EBS   |           |        |
+----+--+----------+----+------+-----+------+---+--------+-----------+--------+
|12  |1 |tp-vlan   | N  |    0 |50048|50048 |16 |       0 |VLN| 101   |extended|
+----+--+----------+----+------+-----+------+---+--------+-----------+--------+
|12  |2 |tp-vlan-1 | N  |    0 |60032|60032 |16 |       0 |VLN| 102   |extended|
+----+--+----------+----+------+-----+------+---+--------+-----------+--------+
|12  |3 |tp-vlan-2 | N  |    0 |70016|70016 |16 |       0 |VLN| 103   |extended|
+----+--+----------+----+------+-----+------+---+--------+-----------+--------+
|12  |4 |tp-vlan-3 | N  |    0 |80000|80000 |16 |       0 |VLN| 104   |extended|
+----+--+----------+----+------+-----+------+---+--------+-----------+--------+
|12  |5 |tp-vlan-4 | N  |    0 |90048|90048 |16 |       0 |VLN| 105   |default |
+----+--+----------+----+------+-----+------+---+--------+-----------+--------+
```


Copyright© 2022 Ciena® Corporation

This example shows sample output for displaying all 51XX 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               |
+------+---+---------------+----+------+-------+-------+-------+---------+------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 8-54
Clearing statistics for all standard profiles

Statistics count the number of accepted and dropped bytes and packets for each standard profile.

> ***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
1	Clear statistics for all standard profiles: `traffic-profiling standard-profile clear statistics` **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 8-55
Clearing statistics for all global profiles

Clear statistics for all global profiles.

Step	Action
1	Clear statistics for all global profiles: `traffic-profiling global-profile clear statistics` **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 8-56
Clearing statistics for all global parent profiles

Clear statistics for all global parent profiles.

Step	Action
1	Clear statistics for all global parent profiles: `traffic-profiling global-profile parent clear statistics`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 8-57
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

This 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© 2022 Ciena® Corporation

Procedure 8-58
Displaying statistics for all global profiles

Display statistics for all global profiles.

Step	Action
1	Display statistics for all global profiles: `traffic-profiling global-profile show statistics [global-profile-object-list]`

—end—

Example

This example shows sample output for displaying statistics for all global profiles.

```
>traffic-pr g show stat

+----------------------------- GLOBAL PROFILE TABLE -------------------------------+
|     | Profile             | Statistics                                            |
|Prec | Name                | Type       Bytes                  Packets             |
+-----+---------------------+-------------------------------------------------------+
|20   | test                | Accepted   12779520               | 49920               |
|     |                     | Dropped    6111078144             | 23871408            |
+-----+---------------------+-------------------------------------------------------+
|30   | test2               | Accepted   0                      | 0                   |
|     |                     | Dropped    0                      | 0                   |
+-----+---------------------+-------------------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 8-59
Displaying statistics for global parent profiles

Display statistics for

- all global parent profiles
- a specific parent profile

Step	Action

To display statistics for all global parent profiles

1 Display statistics for all global parent profiles:

```
traffic-profiling global-profile parent show statistics
```

To display statistics for a specific global parent profile

2 Display statistics for a specific global parent profile:

```
traffic-profiling global-profile parent show statistics
profile <global-profile-object>
```

—end—

Example

This example shows sample output for displaying statistics for all global parent profiles.

```
>traffic-pr g parent show stat

+------------------------ GLOBAL PARENT PROFILE TABLE ------------------------+
| Profile              | Statistics                                            |
| Name                 | Type       Bytes                  Packets              |
+----------------------+-------------------------------------------------------+
| p1                   | Accepted   0                      | 0                  |
|                      | Dropped    0                      | 0                  |
+----------------------+-----------------------------------+-------------------+
```

This example shows sample output for displaying statistics for a specific global parent profile.

```
>traffic-pr g parent show stat p1

+------------------------------ GLOBAL PROFILE TABLE -------------------------------+
|      | Profile              | Statistics                                           |
|Prec  | Name                 | Type       Bytes                  Packets            |
+------+----------------------+-------------------------------------------------------+
|      | p1                   | **** Child Profile List ****                         |
+------+----------------------+-----------------------------------+-------------------+
|10    | ch1                  | Accepted   0                      | 0                 |
|      |                      | Dropped    0                      | 0                 |
+------+----------------------+-----------------------------------+-------------------+
|11    | ch2                  | Accepted   0                      | 0                 |
|      |                      | Dropped    0                      | 0                 |
+------+----------------------+-----------------------------------+-------------------+
|      |                      | **** Child Profiles Sum ****                         |
+------+----------------------+-----------------------------------+-------------------+
|      | p1                   | Accepted   0                      | 0                 |
|      |                      | Dropped    0                      | 0                 |
+------+----------------------+-----------------------------------+-------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 8-60
Displaying attributes for all global profiles

Display attributes for

- all global profiles
- a specific global profile

Step	Action

To display attributes for all global profiles

1 Display attributes for all global profiles:

```
traffic-profiling global-profile show
```

To display attributes for a specific global profile

2 Display attributes for all global profiles:

```
traffic-profiling global-profile show [global-profile-
object-list]
```

—end—

Example

This example shows sample output for displaying attributes for all global profiles.

```
>traffic-pr g show
+------------------------------------------------------------ GLOBAL PROFILE TABLE ------------------------------------------------------------+
|                  |      |Admin|Oper |     Bandwidth(Kbps)     |    Max Burst(KB)    |                    Classifiers                    |
| Profile Name     | Prec |State|State| CIR        | PIR        | CBS     | EBS     | Type         | Value                             |
+------------------+------+-----+-----+------------+------------+---------+---------+--------------+-----------------------------------+
|ch1             *| 10   | Dis | Dis | 1024       | 1024       | 12      | 0       | DSCP         | 10,20                             |
+------------------+------+-----+-----+------------+------------+---------+---------+--------------+-----------------------------------+
|ch2             *| 11   | Dis | Dis | 2048       | 2048       | 12      | 0       | Vlan         | 100,101                           |
+------------------+------+-----+-----+------------+------------+---------+---------+--------------+-----------------------------------+
* denotes a child profile
```


Copyright© 2022 Ciena® Corporation

This example shows sample output for displaying attributes for a specific global profile.

```
> traffic-pr g show profile ch1
+---------------------------- GLOBAL TRAFFIC PROFILE ----------------------------+
| Parameter                | Value                                                |
+--------------------------+------------------------------------------------------+
| Profile Name             | ch1                                                  |
| Precedence               | 10                                                   |
| Admin State              | Disabled                                             |
| Operational State        | Disabled                                             |
| Is Child                 | True                                                 |
| Parent Profile           | p1                                                   |
| CIR (Kbps)               | 1024                                                 |
| PIR (Kbps)               | 1024                                                 |
| CBS (Kbytes)             | 12                                                   |
| EBS (Kbytes)             | 0                                                    |
| DSCP Remark Policy       | leave                                                |
| Fixed DSCP Value         | 0                                                    |
| Stats Collection         | On                                                   |
| ingress-color-aware      | On                                                   |
| remark-color-policy      | leave                                                |
| remark-rcos-policy       | leave                                                |
| yellow-remark-rcos       | 2                                                    |
| green-remark-rcos        | 2                                                    |
+--------------------------+------------------------------------------------------+
|                     CLASSIFIERS                                                 |
+--------------------------+------------------------------------------------------+
| DSCP                     | 10,20                                                |
+--------------------------+------------------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 8-61
Displaying attributes for global parent profiles

You can display attributes for

- a specific global parent profile
- all global parent profiles

Step	Action

To display attributes for a specific global parent profile

1 Display attributes for a specific global parent profiles:

```
traffic-profiling global-profile parent show [<global-
profile-object>]
```

To display attributes for all global parent profiles

2 Display attributes for all global parent profiles:

```
traffic-profiling global-profile parent show
```

—end—

Example

This example displays attributes for the global parent profile p1.

```
> traffic-profiling global-profile parent show profile p1

+----------------------- GLOBAL PARENT TRAFFIC PROFILE ------------------------+
| Parameter              | Value                                                |
+------------------------+------------------------------------------------------+
| Profile Name           | p1                                                   |
| Admin State            | Disabled                                             |
| Num Children           | 2                                                    |
| CIR (Kbps)             | 10048                                                |
| CBS (Kbytes)           | 12                                                   |
+------------------------+------------------------------------------------------+
|                   CLASSIFIERS                                                  |
+------------------------+------------------------------------------------------+
| Port                   | 5                                                    |
+------------------------+------------------------------------------------------+
+-------------------------------------------- GLOBAL CHILD PROFILE TABLE ---------------------------------------------+
|                 |      |Admin|Oper |     Bandwidth(Kbps)     |    Max Burst(KB)    |                Classifiers                 |
| Profile Name    | Prec |State|State| CIR        | PIR        | CBS      | EBS      | Type        | Value                          |
+-----------------+------+-----+-----+------------+------------+----------+----------+-------------+--------------------------------+
| ch1             | 10   | Dis | Dis | 1024       | 1024       | 12       | 0        | DSCP        | 10,20                          |
+-----------------+------+-----+-----+------------+------------+----------+----------+-------------+--------------------------------+
| ch2             | 11   | Dis | Dis | 2048       | 2048       | 12       | 0        | Vlan        | 100,101                        |
+-----------------+------+-----+-----+------------+------------+----------+----------+-------------+--------------------------------+
```

This example displays attributes for all global parent profiles.

```
> traffic-profiling global-profile parent show

+------------------------------------ GLOBAL PARENT PROFILE TABLE ------------------------------------+
|                  |            |            |            |                     Classifiers                |
| Profile Name     | Num Child  | CIR        | CBS        | Type         | Value                              |
+------------------+------------+------------+------------+--------------+------------------------------------+
| p1               | 2          | 10048      | 12         | Port         | 5                                  |
+------------------+------------+------------+------------+--------------+------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 8-62
Displaying throughput statistics for a global profile

Display throughput statistics for a global profile to monitor the throughput traffic profile from the CLI. Throughput statistics are used for troubleshooting.

Step	Action
1	Display throughput statistics for a global profile: `traffic profiling global-profile show throughput <global-profile-object-list>` where <global-profile-object-list> is the global profile that you want to monitor

—end—

Example

This example output for throughput statistics for the global profile “test”

```
> traffic-pr g show throughput profile test
+------------------------ TRAFFIC-PROFILE test THROUGHPUT ------------------------+
| Statistics        | Current Value       | Delta Value         | Rate Mpps & Mbps    |
+-------------------+---------------------+---------------------+---------------------+
| Time              |          0:00:04:22 |       0:00:00:03.60 |                     |
+-------------------+---------------------+---------------------+---------------------+
| Accepted Bytes    |            14622976 |              425216 |               0.945 |
| Dropped Bytes     |          7009427968 |           207309056 |             460.687 |
| Accepted Pkts     |               57122 |                1662 |               0.000 |
| Dropped Pkts      |            27380599 |              809813 |               0.225 |
+-------------------+---------------------+---------------------+---------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 8-63
Displaying throughput statistics for a global parent profile

Display throughput statistics for a global parent profile to monitor the throughput traffic profile from the CLI. Throughput statistics are used for troubleshooting.

Step	Action
1	Display throughput statistics for a global parent profile: `traffic profiling global-profile parent show throughput <global-profile-object-list>` where <global-profile-object-list> is the global profile that you want to monitor

—end—

Example

This example output for throughput statistics for the global parent profile “p1”

```
> traffic-pr g show throughput parent profile p1
+------------------------- TRAFFIC-PROFILE p1 THROUGHPUT ----------------+
| Statistics        | Current Value    | Delta Value      | Rate Mpps & Mbps |
+-------------------+------------------+------------------+------------------+
| Time              |       0:00:10:08 |   0:00:00:05.47  |                  |
+-------------------+------------------+------------------+------------------+
| Accepted Bytes    |                0 |                0 |                  |
| Dropped Bytes     |                0 |                0 |                  |
| Accepted Pkts     |                0 |                0 |                  |
| Dropped Pkts      |                0 |                0 |                  |
+-------------------+------------------+------------------+------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 8-64
Displaying throughput statistics for a standard traffic profile

Display throughput statistics for a standard traffic profile to monitor the throughput 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 <TrafficProfileName> is the traffic profile that you want to monitor

—end—

Copyright© 2022 Ciena® Corporation

Procedure 8-65
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© 2022 Ciena® Corporation

Procedure 8-66
Setting the child mode

Set the child-mode attribute to determine the classification of child traffic profiles. This table lists valid values for the child-mode attribute.

Table 8-40
Valid values for the child-mode attribute

Value	Description
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 on the VLAN ID.
vlan-cos	Finds a matching profiling based on 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© 2022 Ciena® Corporation

Procedure 8-67
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

This 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© 2022 Ciena® Corporation

Procedure 8-68
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

This 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© 2022 Ciena® Corporation

Procedure 8-69
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

This 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© 2022 Ciena® Corporation

Procedure 8-70
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

This 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© 2022 Ciena® Corporation

Procedure 8-71
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

This 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© 2022 Ciena® Corporation

Procedure 8-72
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 <Congestion AvoidanceProfile>	identifies the specific congestion avoidance profile.

—end—

Example

This 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)    |
+-----------+-----------------------++-----------+-------------------------+
```


Copyright© 2022 Ciena® Corporation

This example shows the default sWRED congestion avoidance profile on the 5160.

```
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© 2022 Ciena® Corporation

Procedure 8-73
Updating the congestion avoidance profile for an egress port queue

The default congestion avoidance profile for egress port queues is to use Default-S-WRED. 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

This 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© 2022 Ciena® Corporation

Procedure 8-74
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

This 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© 2022 Ciena® Corporation

Procedure 8-75
Displaying the congestion avoidance profile for an egress port queue

Display the congestion avoidance profile for an egress port queue.

Step	Action
1	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

This example shows sample output for the Default-S-WRED congestion avoidance profile on a 5160 switch.

```
traffic-services queuing egress-port-queue-group show queue 0 port 1.1

+------------------------------- 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     |
+---------+----------+---------+----------+----------+----------+---------+
```


Copyright© 2022 Ciena® Corporation

Procedure 8-76
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 <Congestion AvoidanceProfile> identifies the congestion avoidance profile to be deleted.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 8-77
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 <Congestion AvoidanceProfile> identifies the congestion avoidance profile to be renamed. name <String[15]> is the new name for the congestion avoidance profile. **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 8-78
Changing the algorithm of an egress port scheduler for non-hierarchical egress queuing

Change the algorithm of an egress port scheduler for non-hierarchical egress queuing.

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

This 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© 2022 Ciena® Corporation

Procedure 8-79
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 must 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>` ***Note:*** For 3924, 3926, 3928, 3942, 5142 and 5160 devices only, the scheduler-weight range is 0, 2-127. **—end—**

Example

This example changes a queue's scheduler weight.

```
traffic-services queuing egress-port-queue-group set queue 0 port 9 scheduler-
weight 100
```


Copyright© 2022 Ciena® Corporation

Procedure 8-80
Displaying queue weight and scheduler algorithms for non-hierarchical egress queuing

Display queue weight and scheduler algorithms for non-hierarchical egress queuing.

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

This example displays queue weight and scheduler algorithms.

```
traffic-services queuing egress-port-queue-group show port 9

+----------------------------  QUEUE GROUPDATA ----------------------------+
|                                                                           |
+-------------------------------+---------------------------------------+
| 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© 2022 Ciena® Corporation

Procedure 8-81
Configuring egress port and queue shaping for non-hierarchical egress queuing

You can set

- egress port shaper-rate bandwidth
- egress port shaper-rate burst-size
- egress shaping parameter for each queue

This table lists the steps that the shaper rate increments in based on the burst size.

Table 8-41
Burst size to shaper-rate increment

Burst size	Shaper rate increments in steps of
0 or less than 16 M	64 kbps
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

This table lists the steps that the CIR and PIR increments in based on the shaper rate.

Table 8-42
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
16.78 Gbps and less than 33.55 Gbps	128 kbps
33.55 Gbps and less than 40 Gbps	256 kbps

Copyright© 2022 Ciena® Corporation

This table lists the parameters for individual shaping of queues.

Table 8-43
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 3903, 3903x, 3904, 3905, and 3906 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..80000000>
```

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..262144>
```


Copyright© 2022 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..80000000> cbs <NUMBER: 0..262144> eir <NUMBER:
0..80000000> ebs <NUMBER: 0..262144>
```

—end—

Example

This example sets the egress port shaper-rate bandwidth.

```
traffic-services queuing egress-port-queue-group set port 3 shaper-rate 1200
```

This example sets the egress port shaper-rate burst-size.

```
traffic-services queuing egress-port-queue-group set port 3 burst-size 1024
```

This 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© 2022 Ciena® Corporation

Procedure 8-82
Displaying egress port queue configuration for non-hierarchical egress queuing

You can display egress port queue configuration for non-hierarchical egress queuing:

- 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—

Copyright© 2022 Ciena® Corporation

Example

This example displays the output for the traffic-services queuing egress-port-queue-group show port command.

```
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                                   |
| Scheduling Algorithm           | strict                              |
| Shaper Bandwidth (Kbps)        | 20000000                            |
| Shaper Burst Size (Kb)         | 10240                               |
+------------------------------------------------------------------------+
|    |      Scheduler     |  Size       CIR    |   CBS     |   EIR    |   EBS    |
| Q  | Pri Idx | Weight   | (Pckts) | (Kbps)   | (Kbytes)  | (Kbps)   | (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© 2022 Ciena® Corporation

Procedure 8-83
Displaying egress port queue statistics for non-hierarchical egress queuing

Statistics count the bytes and packets dropped and transmitted by the egress port queues.

You can display egress port queue statistics 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
```

where

queue <NUMBER: 0..7>	is the queue in the queue group
port <PortQueueGroup>	is the port on which the queue group is instantiated.

Copyright© 2022 Ciena® Corporation

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© 2022 Ciena® Corporation

Procedure 8-84
Configuring hierarchical egress port shaping

Configure hierarchical egress port shaping.

Note: This procedure applies to 3924, 3926, 3928, 3942, 5142 and 5160 devices only.

Step	Action
1	Set the egress port shaper-rate bandwidth: `traffic-services queuing egress-port <port-name> shaper-rate <NUMBER: 0..80000000>` where <port-name> — is the port you are configuring hierarchical egress port shaping on. shaper-rate <NUMBER: 0..80000000> — is the shaper rate.
2	Set the egress port burst size: `traffic-services queuing egress-port set port <port-name> burst-size <NUMBER: 4..262144>` where <port-name> — is the port you are setting the burst size on. burst-size <NUMBER: 4..262144> — is the burst size.

—end—

Example

This example configures hierarchical egress port shaping.

```
traffic-services queuing egress-port 5 shaper-rate 500000
traffic-services queueing egress-port set port 5 burst-size 1024
```


Copyright© 2022 Ciena® Corporation

Procedure 8-85
Configuring the default queue group for hierarchical egress queuing

Configure the default queue group for hierarchical egress queuing.

Note: This procedure applies to 3924, 3926, 3928, 3942, 5142 and 5160 devices only.

Step	Action
1	Set the rate parameters of the default queue group: `traffic-services queuing egress-port-queue-group set port <port-name> cir <NUMBER: 1024..80000000> cbs <NUMBER: 0..262144> eir <NUMBER: 0..80000000> scheduler-weight <0..127> eir-policy <configure | inherit>` where <port-name> — is the port you are configuring the default queue group. cir <NUMBER: 1024..80000000> — sets the CIR. cbs <NUMBER: 0..262144> — sets the CBS. eir <NUMBER: 0..80000000> — sets the EIR. scheduler-weight <0..127> — sets the scheduler weight. ***Note:*** For 3924, 3926, 3928, 3942, 5142 and 5160 devices only, the scheduler-weight range is 0, 2-127. eir-policy <configure | inherit> — sets the EIR policy to configure or inherit.

Copyright© 2022 Ciena® Corporation

2 Set the parameters of the queue of the default queue group:

```
traffic-services queuing egress-port-queue-group set
queue <NUMBER: 0..7> port <port-name> cir <NUMBER:
0..80000000> cbs <NUMBER: 0..262144> eir <NUMBER:
0..80000000> scheduler-weight <0..127> eir-policy
<configure | inherit> congestion-avoidance-profile
<profile-name>
```

where

queue <NUMBER: 0..7>	sets the queue for the default queue group.
<port-name>	is the port for the default queue group.
cir <NUMBER: 1024..80000000>	sets the CIR queue of default queue group.
cbs <NUMBER: 0..262144>	sets the CBS queue of the default queue group.
eir <NUMBER: 0..80000000>	sets the EIR queue of the default queue group.
scheduler-weight <0, 10-127>	sets the scheduler weight queue of the default queue group. ***Note:*** For 3924, 3926, 3928, 3942, 5142 and 5160 devices only, the scheduler-weight range is 0, 2-127.
eir-policy <configure | inherit>	sets the EIR policy to configure or inherit.
congestion-avoidance-profile <profile-name>	sets the congestion-avoidance-profile queue of the default queue group.

—end—

Example

This example configures the default queue group.

```
traffic-services queuing egress-port-queue-group set port 5 cir 500000 cbs
262000 eir 500000 scheduler-weight 125 eir-policy inherit
traffic-services queuing egress-port-queue-group set queue 5 port 5 cir 500000
cbs 262000 eir 500000 scheduler-weight 125 eir-policy inherit congestion-
avoidence-profile profile1
```


Copyright© 2022 Ciena® Corporation

Procedure 8-86
Configuring the queue group profile for hierarchical egress queuing

Configure the queue group profile for hierarchical egress queuing.

Note: This procedure applies to 3924, 3926, 3928, 3942, 5142 and 5160 devices only.

Step	Action
1	Create a queue group profile or dynamic queue group profile: `traffic-services queuing egress-queue-group-profile create <profile-name>` where profile <profile-name> — is the queue group profile name.
2	Set the parameters of the queue of the queue group profile: `traffic-services queuing egress-queue-group-profile set queue <NUMBER: 0..7> profile <profile-name> cir <NUMBER: 0..80000000> cbs <NUMBER: 0..262144> eir <NUMBER: 0..80000000> ebs <NUMBER: 0..262144. scheduler-weight <0..127> eir-policy <configure | inherit> congestion-avoidance-profile <congestion profile-name>` where queue <NUMBER: 0..7> — sets the queue for the queue group profile. profile <profile-name> — is the name of the queue group profile. cir <NUMBER: 1024..80000000> — sets the cir for the specified queue for the queue group profile. cbs <NUMBER: 0..262144> — sets the cbs for the specified queue of the queue group profile. eir <NUMBER: 0..80000000> — sets the eir for the specified queue of the queue group profile. ebs <NUMBER: 0..262144> — sets the ebs for the specified queue of the queue group profile.

Copyright© 2022 Ciena® Corporation

where

scheduler-weight <0..127>	sets the scheduler weight for the specified queue of the queue group profile.
eir-policy <configure | inherit>	sets the eir policy for the specified queue of the queue group profile.
congestion-avoidance-profile <profile-name>	sets the congestion-avoidance-profile for the specified queue of the queue group profile.

—end—

Example

This example configures the queue group profile:

```
traffic-services queuing egress-queue-group-profile create profile profile1
traffic-services queuing egress-queue-group profile set queue 5 profile
profile1 cir 50000 cbs 1024 eir 50000 scheduler-weight 125 eir-policy inherit
congestion-avoidence-profile profile1
```


Copyright© 2022 Ciena® Corporation

Procedure 8-87
Allocating the number of queue groups for each physical port on the 3924, 3926 and 3928

Due to hardware limitations on the 3924, 3926 and 3928, the ports pre-allocate the queue groups that are required for the device.

Note: This procedure must be done before any queue group is created.

Step	Action
1	Allocate the number of queue groups for each physical port: `traffic-services queuing resource set port <port num> queue-group-count <0..7>` where port <port num> — is the physical port you are allocating the number of queue groups to. queue-group-count <0..7> — is the queue-group count for the physical port.

—end—

Example

This example allocates the number of queue groups for each physical port:

```
> traffic-services queuing set resource port 1 queue-group-count 3
> traffic-services queuing set resource port 2 queue-group-count 4
```

Once the queue groups are allocated, the queue group is created:

```
> traffic-services queuing egress-queue-group create name test profile
Default-profile
```

Add ports to the queue group, which must be at the same priority index:

```
> traffic-services queueing egress-queue-group add port 1 queue-group test
prioiryt 1 scheduler-weight 10 cir 1000 cbs 10

> traffic-services queuing egress-queue-group add port 2 queue-group test
priority 1 scheduler-weight 10 cir 1000 cbs 10
```

Setup the VLAN:

```
> vlan create vlan 100
> vlan add vlan 100 port 1, 2
```


Copyright© 2022 Ciena® Corporation

Set up service queuing on VLAN:

```
> vlan queuing set vlan 100 queue-group test
```


Copyright© 2022 Ciena® Corporation

Procedure 8-88
Deallocating the number of queue groups for each physical port on the 3924, 3926 and 3928

You can deallocate the number of queue groups for each physical port on the 3924, 3926 and the 3928.

Step	Action
1	De-allocate the number of queue groups for each physical port: `traffic-services queuing resource unset port <port num> queue-group-count` where port <port num> is the physical port you are de-allocating the number of queue groups from. **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 8-89
Displaying the number of queue groups for each physical port on the 3924, 3926 and 3928

After you allocate the number of queue groups for each physical port on the 3924, 3926 and 3928, you can verify the changes.

Step	Action
1	Display the number of queue groups allocated for each physical port on the 3924, 3926 and 3928: `traffic-services queuing resource show`

—end—

Example

This example shows the output from the traffic-services queuing resource show command.

```
> traffic-services queuing resource show

+----------------------------- QUEUE RESOURCE DATA -------------------------------+
|                                                                                  |
+----------------------------------------------------------------------------------+
| Custom Queues Not Allocated      | 200                                           |
| Custom Queues Allocated          | 56                                            |
| Queue Groups Allocated           | 7                                             |
| Custom Groups Created            | 1                                             |
+----------------------------------------------------------------------------------+
| Port                             | Custom Group Instances | Queue Groups Allocated|
+----------------------------------------------------------------------------------+
| 1                                | 1                      | 3                     |
| 2                                | 1                      | 4                     |
| 3                                | 0                      | 0                     |
| 4                                | 0                      | 0                     |
| 5                                | 0                      | 0                     |
| 6                                | 0                      | 0                     |
| 7                                | 0                      | 0                     |
| 8                                | 0                      | 0                     |
| 9                                | 0                      | 0                     |
| 10                               | 0                      | 0                     |
| 11                               | 0                      | 0                     |
| 12                               | 0                      | 0                     |
+----------------------------------------------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 8-90
Configuring the dynamic queue group profile for hierarchical egress queuing

Configure the dynamic queue group profile for hierarchical egress queuing.

Note: This procedure applies to 3924, 3926, 3928, 3942, 5142 and 5160 devices only.

Step	Action
1	Create a dynamic egress queue group profile: `traffic-services queuing dynamic-egress-queue-group-profile create profile <profile-name>` where profile <profile-name> is the queue group profile name.

Copyright© 2022 Ciena® Corporation

2 Set the parameters of the queue for the dynamic egress queue group profile:

```
traffic-services queuing dynamic-egress-queue-group-
profile set queue <NUMBER: 0..7> profile <profile-name>
cir-percent <NUMBER: 0-100> cbs <NUMBER: 0-13104> eir-
percent <NUMBER: 0-100> ebs <NUMBER: 0-13104> scheduler-
weight <0-127> congestion-avoidance-profile <Congestion
Avoidance Profile Name>
```

where

queue <NUMBER: 0..7>	is the queue.
profile <profile-name>	is the dynamic egress queue group profile.
cir-percent <NUMBER: 0-100>	sets the CIR in percentage for the specified queue in the queue group profile.
cbs <NUMBER: 0-13104>	sets the CBS for the specified queue of the queue group profile.
eir <NUMBER: 0-100>	sets the EIR in percentage for the specified queue of the queue group profile.
ebs <NUMBER: 0-13104>	sets the EBS for the specified queue of the queue group profile.
scheduler-weight <0..127>	sets the scheduler weight for the specified queue of the queue group profile. ***Note:*** For 3924, 3926, 3928, 3942, 5142 and 5160 devices only, the scheduler-weight range is 0, 10-127.
congestion-avoidance-profile <Congestion Avoidance Profile Name>	sets the congestion-avoidance-profile for the specified queue of the queue group profile.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 8-91
Configuring the custom queue group for hierarchical egress queuing

Configure the custom queue group for hierarchical egress queuing.

Note: This procedure applies to 3924, 3926, 3928, 3942, 5142 and 5160 devices only.

Step	Action
1	Create the custom queue group: `traffic-services queuing egress-queue-group create queue-group <name>` where queue-group <name> — is the name of the queue group.
2	Set the profile for the custom queue group: `traffic-services queuing egress-queue-group set queue-group <name> profile <profile-name>` where queue-group <name> — is the name of the queue group. profile <profile-name> — is the name of the profile for the custom queue group.

—end—

Example

This example configures the custom queue group:

```
traffic-services queuing egress-queue-group create queue group queue1
traffic-services queuing egress-queue-group set queue-group queue1 profile
profile1
```


Copyright© 2022 Ciena® Corporation

Procedure 8-92
Adding a port to a custom queue group

Add a port to a custom queue group.

Note: This procedure applies to 3924, 3926, 3928, 3942, 5142 and 5160 devices only.

Step	Action
1	Add a port to a custom queue group: `traffic-services queuing egress-queue-group add port <port-name> queue-group <group-name> priority <NUMBER: 0..63> scheduler-weight <NUMBER: 10..127> cir <NUMBER: 0..80000000. cbs <NUMBER: 0..262144> eir <NUMBER: 0..80000000> ebs <NUMBER: 0..262144> eir-policy <configure | inherit>` where port <port-name> — is the port name. queue-group <group-name> — is the queue group's name. priority <NUMBER: 0..63> — is the priority number. scheduler-weight <NUMBER: 10..127> — is the scheduler weight. cir <NUMBER: 0..80000000> — sets the CIR. cbs <NUMBER: 0..262144> — sets the CBS. eir <NUMBER: 0..80000000> — sets the EIR. ebs <NUMBER: 0..80000000> — sets the EBS. eir-policy <configure | inherit> — sets the EIR policy to configure or inherit.

Copyright© 2022 Ciena® Corporation

2 Set the attributes of the custom queue group for that port:

```
traffic-services queuing egress-queue-group set port
<port-name> queue-group <group-name> priority <NUMBER:
0..63> scheduler-weight <NUMBER: 10..127> cir <NUMBER:
0..80000000. cbs <NUMBER: 0..262144> eir <NUMBER:
0..80000000> ebs <NUMBER: 0..262144> eir-policy
<configure | inherit>
```

where

port <port-name>	is the port name.
queue-group <group-name>	is the queue group's name.
priority <NUMBER: 0..63>	is the priority number.
scheduler-weight <NUMBER: 10..127>	is the scheduler weight.
cir <NUMBER: 0..80000000>	sets the CIR.
cbs <NUMBER: 0..262144>	sets the CBS.
eir <NUMBER: 0..80000000>	sets the EIR.
ebs <NUMBER: 0..80000000>	sets the EBS.
eir-policy <configure | inherit>	sets the EIR policy to configure or inherit.

—end—

Example

This example adds a port to a custom queue group.

```
traffic-services queuing egress-queue-group add port 5 queue-group queue1
priority 50 scheduler-weight 125
traffic-services queuing egress-queue-group set port 5 queue-group group1
priority 50 scheduler-weight 125 cir 500000 cbs 1024 eir 500000 ebs 1024 eir-
policy inherit
```


Copyright© 2022 Ciena® Corporation

Procedure 8-93
Associating a custom queue group with a virtual circuit for hierarchical egress queuing

Note 1: This procedure applies to 3942, 5142 and 5160 devices only.

Note 2: When a physical port is added to an aggregation with an associated custom queue group, the physical port has the same associated custom queue groups with the same settings as the aggregation. The physical port cannot have any associations to custom queue groups for it to be allowed to be added to the aggregation.

To associate a virtual switch with the custom queue group you must

- create a custom queue group and set the queue attributes
- add ports to the custom queue group
- create the virtual switch
- associate the virtual switch with the custom queue group

Traffic disruption may occur when

- configuring a scheduler change
- configuring a priority change
- adding a port to a custom queue group
- setting any attributes for the virtual circuit/VLAN service queuing configuration

Step	Action
Create a custom queue group and set queue attributes	

1	Create an egress queue group profile or dynamic egress queue group profile:
	`traffic-services queuing egress-queue-group-profile create profile <String[1-15]>`
	`traffic-services queuing dynamic-egress-queue-group-profile create profile <profile>`

2	Set queue attributes for the queue-group profile: `traffic-services queuing egress-queue-group-profile set profile <String[1-15]> queue <NUMBER:0-7> cir <NUMBER:0-80000000> eir <NUMBER:0-80000000> cbs <NUMBER: 0-131040>`

Copyright© 2022 Ciena® Corporation

```
ebs <NUMBER:0-131040> scheduler-weight <NUMBER:0-127>
eir-policy<inherit|configure> congestion-avoidance-
profile <STRING>
```

```
traffic-services queuing dynamic-egress-queue-group-
profile set queue <NUMBER: 0..7> profile <profile-name>
cir-percent <NUMBER: 0-100> cbs <NUMBER: 0-13104> eir-
percent <NUMBER: 0-100> ebs <NUMBER: 0-13104> scheduler-
weight <0-127> congestion-avoidance-profile <Congestion
Avoidance Profile Name>
```

3 Create the custom queue-group:

```
traffic-services queuing egress-queue-group create queue-
group <String[1-15]> profile <String[1-15]>
```

where

queue-group <String[1-15]>	is the queue-group name.
profile <String[1-15]>	is the associated queue group profile which provides the attributes of the queue.

Add ports to the custom queue-group

4 Add ports to the custom queue-group:

```
traffic-services queuing egress-queue-group add port
<NUMBER: 1-24> queue-group <custom queue group> priority
<NUMBER: 0-63> {[cir <NUMBER: 0-80000000>] [eir <NUMBER:
0-80000000>][ebs <NUMBER: 0-131040>][cbs <NUMBER: 0-
131040> scheduler-weight <NUMBER: 1-127> eir-policy
<inherit|configure>}
```

where

port <NUMBER: 1-24>	is the port being added to the custom queue-group.
queue-group <custom queue group>	is the custom queue group.
priority <NUMBER: 0-63>	is the priority of the queue group.
cir <NUMBER: 0-80000000>	is the cir value.
eir <NUMBER: 0-80000000>	is the eir value.
ebs <NUMBER: 0-131040>	is the ebs value.

Copyright© 2022 Ciena® Corporation

where

cbs <NUMBER: 0-131040>	is the cbs value.
scheduler-weight <NUMBER: 1-127>	is the weighted value given to the scheduler.
eir-policy <inherit|configure>	determines whether eir policy is inherited from the port.

Create a virtual switch

5 Create a virtual switch:

```
vlan create vlan <VLAN ID List>
```

where

vlan <VLAN ID List>	is the VLAN being created.

6 Add the VLAN to the port:

```
vlan add vlan <VLAN list> port <PortList>
```

where

vlan <VLAN ID List>	is the VLAN being created.
port <PortList>	is a list of ports to add to the VLAN.

7 Add the virtual switch to the reserved VLAN:

```
virtual-switch add reserved-VLAN <VLAN ID List>
```

where

reserved-VLAN <VLAN ID List>	is the reserved VLAN that the virtual switch is being added to.

8 Create an Ethernet virtual circuit:

```
virtual-circuit ethernet create vc <vc> vlan <VLAN>
```

where

vc <vc>	is the virtual circuit name.
vlan <VLAN>	is the provider VLAN.

9 Create an Ethernet virtual switch:

```
virtual-switch ethernet create vs <vs> vc <vc> reserved-
vlan <VLAN ID List>
```

where

vs <vs	is the virtual switch name.
vc <vc>	is the Ethernet virtual circuit.
reserved-vlan <VLAN ID List>	is the reserved VLAN that the virtual switch is being added to.

Copyright© 2022 Ciena® Corporation

10 Add the Ethernet virtual switch to the virtual switch and ports:

```
virtual-switch ethernet add vs <Virtual Switch Name> port
<portlist> vlan <VLAN ID list>
```

where

vs <Virtual Switch Name>	is the virtual switch name.
port <,portlist>	is a list of ports to added to the virtual switch.
vlan <VLAN ID List>	is the VLAN ID list.

Associate a virtual switch with the custom queue group

11 Associate the virtual switch with the custom queue group:

```
virtual-switch Ethernet queuing set vs <Virtual Switch
Name> queue-group <Egress Queue Group> queue-map <queue-
map>
```

where

vs <Virtual Switch Name>	is the virtual switch name.
queue-group <Egress Queue Group>	is the egress queue group.
queue-map <queue-map>	is the associated queue map.

—end—

Example

This example associates a virtual switch with the egress-queue group:

```
traffic-services queuing egress-queue-group-profile create profile test
traffic-services queuing egress-queue-group-profile set queue 3 profile test
cir 10000 cbs 256
traffic-services queuing egress-queue-group create queue-group s1 profile
test

traffic-services queing egress-queue-group add port 2 queue-group s1 priority
15 cir 20000 cbs 256 scheduler-weight 30

vlan create vlan 100
vlan add vlan 100 port 2
virtual-switch add reserved-vlan 4000
virtual-circuit ethernet create vc vc100 vlan 100
virtual-switch ethernet create vs vs100 port 1 vlan 1001

virtual-switch Ethernet queuing set vs vs100 queue-group s1 queue-map Default-
RCOS
```


Copyright© 2022 Ciena® Corporation

Procedure 8-94
Associating a custom queue group with a VLAN for hierarchical egress queuing

Note 1: This procedure applies to 3924, 3926, 3928, 3942, 5142 and 5160 devices only.

Note 2: When a physical port is added to an aggregation with an associated EVC queue group, the physical port has the same associated custom queue groups with the same settings as the aggregation. The physical port cannot have any associations to custom queue groups for it to be allowed to be added to the aggregation.

To associate a VLAN with the custom queue group you must

- create a custom queue group and set the queue attributes
- add ports to the custom queue group
- create the VLAN
- associate the VLAN with the custom queue group

Traffic disruption may occur when

- configuring a scheduler change
- configuring a priority change
- adding a port to a custom queue group
- setting any attributes for the virtual circuit/VLAN service queuing configuration

Step	Action

Create a custom queue group and set queue attributes

1 Create an egress port queue group profile or dynamic egress queue group profile:

```
traffic-services queuing egress-queue-group-profile
create profile <String[1..15]>
```

```
traffic-services queuing dynamic-egress-queue-group-
profile create profile <profile>
```


Copyright© 2022 Ciena® Corporation

2 Set queue attributes for the queue-group profile:

```
traffic-services queuing egress-queue-group profile set
profile <String[1..15]> queue <NUMBER: 0..7> profile test
cir <NUMBER: 0..80000000> cbs <NUMBER: 0..130140> eir
<NUMBER: 0..80000000> ebs <NUMBER: 0..130140> eir-policy
<inherit|configure>
```

```
traffic-services queuing dynamic-egress-queue-group-
profile set queue <NUMBER: 0..7> profile <profile-name>
cir-percent <NUMBER: 0-100> cbs <NUMBER: 0-13104> eir-
percent <NUMBER: 0-100> ebs <NUMBER: 0-13104> scheduler-
weight <0-127> congestion-avoidance-profile <Congestion
Avoidance Profile Name>
```

Note: For 3924, 3926, 3928, 3942, 5142, 5160 devices only, the scheduler-weight range is 0, 10-127.

3 Create the custom queue-group:

```
traffic-services queuing egress-queue-group create queue-
group <String[1..15]> profile <String[1..15]>
```

where

queue-group <String[1..15]>	is the queue-group name.
profile <String[1..15]>	is the associated queue group profile which provides the attributes of the queue.

To add ports to the custom queue-group

4 Add ports to the custom queue-group:

```
traffic-services queuing egress-queue-group add port
<NUMBER: 1..24> queue-group <custom queue group> priority
<NUMBER: 0..63> {[cir <NUMBER: 0..80000000>] [eir
<NUMBER: 0..80000000>][ebs <NUMBER: 0..131040>][cbs
<NUMBER: 0..131040> scheduler-weight <NUMBER: 1..127>}
```

where

port <NUMBER: 1..24>	is the port being added to the custom queue-group.
queue-group <custom queue group>	is the custom queue group.
priority <NUMBER: 0..63>	is the priority given to the port.
cir <NUMBER: 0..80000000>	is the cir value.
eir <NUMBER: 0..80000000>	is the eir value.

Copyright© 2022 Ciena® Corporation

where

ebs <NUMBER: 0..131040>	is the ebs value.
cbs <NUMBER: 0..131040>	is the cbs value.
scheduler-weight <NUMBER: 1..127>	is the weighted value given to the scheduler.

To create a VLAN

5 Create a VLAN

```
vlan create vlan <VLAN ID List>
```

where

vlan <VLAN ID List>	is the VLAN being created.

6 Add the VLAN to the port:

```
vlan add vlan <VLAN list> port <PortList>
```

where

vlan <VLAN ID List>	is the VLAN being created.
port <PortList>	is a list of ports to add to the VLAN.

To associate a VLAN with the custom queue group

7 Associate the VLAN with the custom queue group:

```
vlan queuing set vlan <String{1..4094] queue-group
<Egress Queue Group> queue-map <queue map>
```

where

vs <String[1..4094]	is the virtual switch name.
queue-group <Egress Queue Group>	is the egress queue group.
queue-map <queue map>	is the associated queue map.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example associates a vlan with the egress-queue group profile:

```
traffic-services queuing egress-queue-group-profile create profile test
traffic-services queuing egress-queue-group-profile set queue 3 profile test
cir 10000 cbs 256
traffic-services queuing egress-queue-group create queue-group s2 profile
test

traffic-services queing egress-queue-group add port 2 queue-group s2 priority
15 cir 20000 cbs 256 scheduler-weight 30
traffic-services queing egress-queue-group add port 1 queue-group s2 priority
15 cir 20000 cbs 256 scheduler-weight 40

vlan create vlan 100
vlan add vlan 100 port 1,2

vlan queuing set vlan 100 queue-group s2 queue-map Default-RCOS
```


Copyright© 2022 Ciena® Corporation

Procedure 8-95
Configuring the egress queue depth for hierarchical egress queuing

The egress port queue/egress port can be configured to have a maximum dynamic depth of:

- low
- medium (default)
- medium-high
- high

Step	Action
1	Configure the egress queue depth: `traffic-services queuing egress-port set port <PortQueueGroup> dynamic-max-queue depth <low | medium | medium-high | high>` where port <PortQueueGroup> is the egress port. **—end—**

Example

This command sets the dynamic maximum queue depth on port 1 to high:

```
traffic-services queueing egress-port-queue-group set port 1 dynamic-max-
queue-depth high
```


Copyright© 2022 Ciena® Corporation

Procedure 8-96
Configuring the port scheduler for hierarchical egress

The egress port specifies the scheduler algorithm that applies to the queue groups and the queues. The scheduler is defined by the scheduler algorithm and the associated weight of the scheduled object.

> ***Note:*** The port scheduler cannot be configured with a scheduler algorithm of strict. To achieve strict, set the weight of the associated scheduled object to zero.

Step	Action
1	Configure scheduling on an egress port: `traffic-services queuing egress-port set port <port> scheduler-algorithm <<round-robin | weighteddeficit-round-robin | weighted-round-robin>` where port <port> — is the egress port scheduler-algorithm <<round-robin | weighted deficit-round-robin | weighted-round-robin> — sets the scheduler algorithm for all queue groups attached to the egress port. For details, refer to the table "Scheduling algorithms" on page 8-51. **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 8-97
Displaying hierarchical egress queuing information

You can display hierarchical egress queuing information.

Step	Action

To display the default queue group

1 Display the default queue group:

```
traffic-services queuing egress-port-queue-group show
port <port>
```

where

port <port> is the port on which the queue group is instantiated.

To display all egress queue group profiles

2 Display all egress queue group profiles:

```
traffic-services queuing egress-queue-group-profile show
```

To display a specified egress queue group profile

3 Display a specified egress queue group profile:

```
traffic-services queuing egress-queue-group-profile show
profile <profile>
```

To display all dynamic egress queue group profiles

4 Display all dynamic egress queue group profiles:

```
traffic-services queuing dynamic-egress-queue-group-
profile show
```

To display a specified dynamic egress queue group profile

5 Display a specified dynamic egress queue group profile:

```
traffic-services queuing dynamic-egress-queue-group-
profile show profile <profile>
```

To display a specified queue in a specified dynamic egress queue group profile

6 Display a specified queue in a specified dynamic egress queue group profile:

```
traffic-services queuing dynamic-egress-queue-group-
profile show queue <queue> profile <profile>
```

To display a specified egress queue group

7 Display a specified egress queue group:

```
traffic-services queuing egress-queue-group show queue-
group <queue-group>
```


Copyright© 2022 Ciena® Corporation

To display the egress port group

8 Display the egress port:

```
traffic-services queuing egress-port show port <port>
```

where

port <port> is the port you wish to show.

To display egress port statistics for a specified port

9 Display egress port statistics for a specified port:

```
traffic-services queuing egress-port show port <port>
statistics
```

To display queueing resources

10 Display queuing resources:

```
traffic-services queuing resource show
```

To display queue-group statistics for a specified queue-group

11 Display queue-group statistics for a specified queue-group:

```
traffic-services queuing egress-queue-group show queue-
group <queue-group> statistics
```

—end—

Example

This example shows output for the traffic-services queuing egress-port-queue-group show port 1 command.

```
> traffic-services queuing egress-port-queue-group show port 1
+------------------------------------------ QUEUE GROUP DATA ------------------------------------------+
|                                                                                                        |
+------------------------------+-----------------------------------------------------------------------+
| Name [Port]                  | 1                                                                     |
| Id                           | 1                                                                     |
| Queue Count                  | 8                                                                     |
| CIR (Kbps)                   | 1024                                                                  |
| CBS (Kbytes)                 | 256                                                                   |
| Configured EIR (Kbps)        | 1000000                                                               |
| Actual EIR (Kbps)            | 1000000                                                               |
| EBS (Kbytes)                 | 10240                                                                 |
| Group Scheduler Weight       | 0                                                                     |
| Priority                     | 63                                                                    |
| EIR Policy                   | Inherit                                                               |
+------------------------------------------------------------------------------------------------------+
|        |  Scheduler  |    CIR    |   CBS    |Actual EIR |   EBS    |Config EIR |EIR Policy| Congestion
| Q/Pri  |  Weight     |   (Kbps)  | (KBytes) |  (Kbps)   | (KBytes) |   (Kbps)  |          | Profile
+-------+------------+-----------+----------+-----------+----------+-----------+----------+----------------
| 0      | 0           | 0         | 0        | 1000000   | 256      | 1000000   | Inherit  | Default-S-WRED
| 1      | 0           | 0         | 0        | 1000000   | 256      | 1000000   | Inherit  | Default-S-WRED
| 2      | 0           | 0         | 0        | 1000000   | 256      | 1000000   | Inherit  | Default-S-WRED
| 3      | 0           | 0         | 0        | 1000000   | 256      | 1000000   | Inherit  | Default-S-WRED
| 4      | 0           | 0         | 0        | 1000000   | 256      | 1000000   | Inherit  | Default-S-WRED
| 5      | 0           | 0         | 0        | 1000000   | 256      | 1000000   | Inherit  | Default-S-WRED
| 6      | 0           | 0         | 0        | 1000000   | 256      | 1000000   | Inherit  | Default-S-WRED
| 7      | 0           | 1024      | 256      | 1000000   | 256      | 1000000   | Inherit  | Default-S-WRED
+-------+------------+-----------+----------+-----------+----------+-----------+----------+----------------
```


Copyright© 2022 Ciena® Corporation

This example shows output for the traffic-services queuing egress-queue-group show queue-group test command.

```
> traffic-services queuing egress-queue-group show queue-group test

+---------------------------------------------- QUEUE GROUP DATA ---------------------------------------+
|                                                                                                        |
+----------------------+-----------------------------------------------------------------------------+
| Name                 | test                                                                        |
| Profile              | Default-Profile                                                             |
| VLAN associations    | No VLAN association present                                                 |
| VS associations      | No VS association present                                                   |
+----------------------+-----------------------------------------------------------------------------+
| Port                 |   Scheduler   |  CIR     |  CBS     | Conf EIR  | Actual EIR |  EBS     |  EIR     |
|                      | Pri | Weight  | (Kbps)   | (KBytes) | (Kbps)    | (Kbps)     | (KBytes) | Policy   |
+----------------------+-----+---------+----------+----------+-----------+------------+----------+----------+
| 1                    | 1   | 10      | 0        | 0        | 1000000   | 1000000    | 2048     | inherit  |
+----------------------+-----+---------+----------+----------+-----------+------------+----------+----------+
```

This example shows output for the traffic-services queuing egress-queue-group-profile show command for profile s1 command.

```
> traffic-services queuing egress-queue-group-profile show profile s1

+------------------------------- QUEUE GROUP PROFILE DATA ---------------------------------+
|                                                                                           |
+-------------------------------+-------------------------------------------------------+
| Name                          | s1                                                    |
| Id                            | 2                                                     |
| Queue Count                   | 8                                                     |
| QGroups Using this profile    | None                                                  |
+-------------------------------+-------------------------------------------------------+
|        |Scheduler|   CIR    |  CBS     |  EIR*    |  EBS     |  EIR     | Congestion |
| Q/Pri  |  Weight |  (Kbps)  | (KBytes) | (Kbps)   | (KBytes) | Policy   | Profile    |
+--------+---------+----------+----------+----------+----------+----------+------------+
| 0      | 20      | 0        | 0        | 1000000  | 256      | inherit  | Default-S-WRED |
| 1      | 30      | 0        | 0        | 1000000  | 256      | inherit  | Default-S-WRED |
| 2      | 40      | 0        | 0        | 1000000  | 256      | inherit  | Default-S-WRED |
| 3      | 50      | 0        | 0        | 1000000  | 256      | inherit  | Default-S-WRED |
| 4      | 60      | 0        | 0        | 1000000  | 256      | inherit  | Default-S-WRED |
| 5      | 70      | 0        | 0        | 1000000  | 256      | inherit  | Default-S-WRED |
| 6      | 80      | 0        | 0        | 1000000  | 256      | inherit  | Default-S-WRED |
| 7      | 90      | 0        | 0        | 1000000  | 256      | inherit  | Default-S-WRED |
+-------------------------------------------------------------------------------------------+
*EIR valid only if EIR Policy set to configure
```

This example shows output for the traffic-services queuing egress-port show port 1 command.

```
> traffic-services queuing egress-port show port 1

+------------------------------------- EGRESS PORT DATA -------------------------------------+
|                                                                                              |
+--------------------------------+-------------------------------------------------------------+
| Name [Port]                    | 1                                                           |
| Id                             | 1                                                           |
| Queue Group Count              | 2                                                           |
| Scheduling Algorithm           | weighted-round-robin                                        |
| Traffic Shaping                | On                                                          |
| Shaper Bandwidth (Kbps)        | 1000000                                                     |
| Shaper Burst Size (KBytes)     | 10240                                                       |
| Dynamic Max Queue Depth        | medium                                                      |
+--------------------------------+-------------------------------------------------------------+
| Queue Group                      | Scheduler    |    CIR     |   CBS      |    EIR     |   EBS     |
|                                  | Pri | Weight |   (Kbps)   | (KBytes)   |   (Kbps)   | (KBytes)  |
+----------------------------------+-----+--------+------------+------------+------------+-----------+
| Default                          | 63  | 0      | 1024       | 256        | 1000000    | 10240     |
+----------------------------------+-----+--------+------------+------------+------------+-----------+
| test                             | 1   | 10     | 0          | 0          | 1000000    | 2048      |
+----------------------------------+-----+--------+------------+------------+------------+-----------+
```


Copyright© 2022 Ciena® Corporation

This example shows output for the traffic-services queuing resource show command:

```
> traffic-services queuing resource show

+----------------------------- QUEUE RESOURCE DATA -------------------------------+
|                                                                                  |
+----------------------------------------------------------------------------------+
| Total Custom Queues remaining    | 768                                           |
| Total Custom Queues used         | 8                                             |
| Custom Groups Created            | 1                                             |
+----------------------------------------------------------------------------------+
| Port                             | Custom Group Instances | Custom Queues used   |
+----------------------------------------------------------------------------------+
| 1                                | 1                      | 8                    |
| 2                                | 0                      | 0                    |
| 3                                | 0                      | 0                    |
| 4                                | 0                      | 0                    |
| 5                                | 0                      | 0                    |
| 6                                | 0                      | 0                    |
| 7                                | 0                      | 0                    |
| 8                                | 0                      | 0                    |
| 9                                | 0                      | 0                    |
| 10                               | 0                      | 0                    |
| 11                               | 0                      | 0                    |
| 12                               | 0                      | 0                    |
| 13                               | 0                      | 0                    |
| 14                               | 0                      | 0                    |
| 15                               | 0                      | 0                    |
| 16                               | 0                      | 0                    |
| 17                               | 0                      | 0                    |
| 18                               | 0                      | 0                    |
| 19                               | 0                      | 0                    |
| 20                               | 0                      | 0                    |
| 21                               | 0                      | 0                    |
| 22                               | 0                      | 0                    |
| 23                               | 0                      | 0                    |
| 24                               | 0                      | 0                    |
+----------------------------------------------------------------------------------+
```

This example shows output for the traffic-services queuing egress-queue-group show queue-group test statistics command:

```
> traffic-services queuing egress-queue-group show queue-group test
statistics

+------------------------------------ QUEUE GROUP STATISTICS ------------------------------------+
|                                                                                                |
|------------------------------------------------------------------------------------------------|
|                                                                                                |
|  Queue Group : test                                                                            |
|                                                                                                |
|  Port : 1                                                                                      |
|                                                                                                |
|  Q ID |     Dropped Bytes     |     Dropped Pkts      |   Transmitted Bytes   |   Transmitted Pkts   |
|-------+-----------------------+-----------------------+-----------------------+----------------------|
|  0    |                     0 |                     0 |                     0 |                    0 |
|  1    |                     0 |                     0 |                     0 |                    0 |
|  2    |                     0 |                     0 |                     0 |                    0 |
|  3    |                     0 |                     0 |                     0 |                    0 |
|  4    |                     0 |                     0 |                     0 |                    0 |
|  5    |                     0 |                     0 |                     0 |                    0 |
|  6    |                     0 |                     0 |                     0 |                    0 |
|  7    |                     0 |                     0 |                     0 |                    0 |
+------------------------------------------------------------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

This example shows output for the traffic-services queuing egress-port show port 1 statistics command:

```
> traffic-services queuing egress-port show port 1 statistics

+------------------------------------- EGRESS PORT STATISTICS ------------------------------+
|                                                                                             |
|---------------------------------------------------------------------------------------------|
| Port: 1                                                                                     |
| Queue group    |  Dropped Bytes  |  Dropped Pkts  |  Transmitted Bytes |  Transmitted Pkts  |
|----------------+-----------------+----------------+--------------------+--------------------|
| Default        |               0 |              0 |              52474 |                488 |
|----------------+-----------------+----------------+--------------------+--------------------|
| test           |               0 |              0 |                  0 |                  0 |
|----------------+-----------------+----------------+--------------------+--------------------|
```


Copyright© 2022 Ciena® Corporation

Procedure 8-98
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.

Note that when the system is in payload calculation mode, the 20-byte IFG is not considered in the bandwidth equations. The physical port must still operate according to Ethernet standards and take into account the IFG when transmitting packets. The operator can possibly configure a CIR sum at a port that exceeds 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:

- 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)

When the bandwidth calculation mode is set to payload, it can cause problems when meeting all queue groups or port CIRs if the sum total is equal to the port speed.

When the bandwidth calculation mode is payload, it means the shaper rates determine the rate based on only the payload portion of a frame, for example not including the IFG and preamble. However, since the IFG and preamble are still consuming a portion of line bandwidth, it is possible to oversubscribe the line.

Copyright© 2022 Ciena® Corporation

The following table describes bandwidth calculation mode.

Table 8-44
Bandwidth calculation mode

				Bandwidth calculation mode	
				Transport *	**Payload ****
	Group	**Packet Size (B)**	**CIR (Mbps)**	**Actual BW**	**Required BW**
	Default	400	100	100	105
	Grp1	400	300	300	315
	Grp2	400	400	400	420
	Grp3	400	200	200	210
Total			1000	1000	1050

* Transport includes packet size, preamble and IFG. This accounts for the total line BW consumed by this group when shaped at the configured rate.

** Payload includes only the packet size. When preamble and IFG are included, the group consumes a larger portion of the available BW.

Step	Action

To set the bandwidth calculation mode

1 Set the bandwidth calculation mode:

```
flow bw-calculation-mode set mode <payload | transport>
```

To calculate bandwidth without IFG:

2 Calculate bandwidth without IFG:

```
flow bw-calculation set mode payload
```

To use IFG in bandwidth calculations

3 Use IFG in bandwidth calculations:

```
flow bw-calculation set mode transport
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 8-99
Sample configuration for local CPU-generated control, management and OAM traffic QoS treatment

This configuration example shows how to enable and modify a protocol COS map, how to set the PCP control to dscp-mapping, and how to set the CPU port RCoS map.

Enable protocol-cos-map (by default map is enabled with default values)

```
> traffic-services cos-mapping protocol-cos-map enable cos-map DefaultProtoCos
```

Disable protocol-cos-map

```
> traffic-services cos-mapping protocol-cos-map disable cos-map DefaultProtoCos
```

Modify the DSCP (24) of any l3 protocol (ldp)

```
> traffic-services cos-mapping protocol-cos-map set cos-map DefaultProtoCos protocol
ldp dscp-cos 24
```

Modify PCP (5) of any l2 protocol (arp)

```
> traffic-services cos-mapping protocol-cos-map set cos-map DefaultProtoCos protocol
arp dot1d-cos 5
```

Set the DSCP of any l3 protocol (ldp) to default value

```
> traffic-services cos-mapping protocol-cos-map unset cos-map DefaultProtoCos protocol
ldp dscp-cos
```

Set the PCP of any l2 protocol (arp) to default value

```
> traffic-services cos-mapping protocol-cos-map unset cos-map DefaultProtoCos protocol
arp dot1d-cos
```

Modify the PCP control to DSCP-mapping

```
> traffic-services cos-mapping set cpu-l3-protocol-pcp-ctrl dscp-mapping
```

Modify the PCP control to interface-pri (unset command also sets the pcp control to default value)

```
> traffic-services cos-mapping set cpu-l3-protocol-pcp-ctrl interface-pri
> traffic-services cos-mapping unset cpu-l3-protocol-pcp-ctrl
```

Modify the CPU port RCoS map

```
> traffic-services queuing queue-map set rcos-map CpuPort-RCOS rcos 5 queue 5
```


Copyright© 2022 Ciena® Corporation

Set the DSCP mode to preserve

```
> traffic-services cos-mapping protocol-cos-map set cos-map DefaultProtoCos protocol
dhcp-l3-relay dscp-mode preserve
```

Unset the DSCP mode

```
> traffic-services cos-mapping protocol-cos-map unset cos-map DefaultProtoCos protocol
dhcp-l3-relay dscp-mode
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 8-100
Configuring a device for dynamic microflow detection and greedy flow monitoring

Perform the following steps to get this feature working on a device.

Step	Action
1	Allocate the appropriate number of resources to traffic-profiling using resource-manager.
2	Change the traffic-profiling option “profile-mode”: `traffic-profiling set profile-mode global-l2-l3-extended`
3	Set the traffic-profiling option “profile-type” to “dynamic-profile”: `traffic-profiling set profile-type dynamic-profile`
4	Enable traffic-profiling: `traffic-profiling enable`

5	Configure global parameters related to microflow detection and policing:
	`traffic-profiling dynamic-profile set/unset service-flow-monitor-timer <>`
	`traffic-profiling dynamic-profile set/unset microflow-monitor-timer <>`
	`traffic-profiling dynamic-profile set/unset greedy-microflow-monitor-timer <>`
	`traffic-profiling dynamic-profile set/unset microflow-batch-monitor-count <>`

6	Create a profile to define the behavior of microflow detection and greedy microflow identification: `traffic-profiling dynamic-profile create name <profile-name>`

Copyright© 2022 Ciena® Corporation

7 Set profile-related parameters to define the behavior of microflow detection and greedy microflow identification:

```
traffic-profiling dynamic-profile set/unset profile <>
service-flow-cir <>
```

```
traffic-profiling dynamic-profile set/unset profile <>
service-flow-cbs <>
```

```
traffic-profiling dynamic-profile set/unset profile <>
service-flow-utilization-threshold <>
```

```
traffic-profiling dynamic-profile set/unset profile <>
service-flow-over-threshold-count <>
```

```
traffic-profiling dynamic-profile set/unset profile <>
service-flow-under-threshold-count <>
```

```
traffic-profiling dynamic-profile set/unset profile <>
greedy-microflow-cir <>
```

```
traffic-profiling dynamic-profile set/unset profile <>
greedy-microflow-cbs <>
```

```
traffic-profiling dynamic-profile set/unset profile <>
greedy-microflow-drop-threshold <>
```

```
traffic-profiling dynamic-profile set/unset profile <>
greedy-microflow-under-threshold-count <>
```

```
traffic-profiling dynamic-profile set/unset profile <>
microflow-over-threshold-count <>
```

```
traffic-profiling dynamic-profile set/unset profile <>
microflow-under-threshold-count <>
```

```
traffic-profiling dynamic-profile set/unset profile <>
microflow-monitor-count <>
```

8 Associate a service flow with the behavior profile:

```
traffic-profiling dynamic-profile add profile <> port <>
vs <> vlan <>
```

9 Enable the behavior profile:

```
traffic-profiling dynamic-profile enable profile <>
```

10 Enable the dynamic profile:

```
traffic-profiling dynamic-profile enable
```


Copyright© 2022 Ciena® Corporation

11 Use the following CLI commands to display various parameters:

```
traffic-profiling show
traffic-profiling dynamic-profile show
traffic-profiling dynamic-profile show detected-
microflows
traffic-profiling dynamic-profile show greedy-microflows
traffic-profiling dynamic-profile show profile <>
traffic-profiling dynamic-profile show profile <>
detected-microflows
traffic-profiling dynamic-profile show profile <> greedy-
microflows
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 8-101
Minimally configuring microflow policing on port 1

Use this procedure to enable microflow policing on port 1.

Step	Action

1	Enter the following (minimal configuration) CLI commands:
	`traffic-profiling set profile-mode global-l2-l3-extended`
	`traffic-profiling set profile-type dynamic-profile`
	`traffic-profiling enable`
	`traffic-profiling dynamic-profile create name test`
	`traffic-profiling dynamic-profile add profile test port 1`
	`traffic-profiling dynamic-profile enable profile test`
	`traffic-profiling dynamic-profile enable`

2	Verify the configuration by entering the following CLI commands:
	`traffic-profiling show`
	`traffic-profiling dynamic-profile show`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 8-102
Working examples of dynamic microflow detection and greedy flow monitoring

Working examples are added here to describe feature use cases.

Example 1

In this use case, port 1 and port 2 are members of VLAN 10. Port 1 is receiving traffic with the following types of streams:

- IPv4 traffic with source IP 10.10.11.12 and destination IP 10.10.11.23 (Rate 5 Mbps)
- IPv4 traffic with source IP 10.10.11.13 and destination IP 10.10.11.24 (Rate 50 Mbps)
- IPv4 traffic with source IP 10.10.11.14 and destination IP 10.10.11.25 (Rate 60 Mbps)
- IPv4 traffic with source IP 10.10.11.15 and destination IP 10.10.11.26 (Rate 80 Mbps)

The user requirements are:

- Traffic ingressing on Port 1 with VLAN 10 must not consume more than 200Mbps of the port's total bandwidth.
- Once traffic with VLAN 10 consumes more than 100 Mbps of allocated bandwidth, then the device must limit the identified/detected greedy microflows at a 10 Mbps rate.
- Greedy microflows must be detected as fast as possible, with some cool off time for momentary threshold crosses.
- All detected microflows must be monitored 10 times.
- Greedy microflow must be metered until its drop rate becomes less than 1 Mbps.

Step	Action
1	The configuration needed to achieve this (on a 5142 device) is:

```
traffic-profiling set profile-mode global-l2-l3-extended

traffic-profiling set profile-type dynamic-profile

traffic-profiling enable

traffic-profiling dynamic-profile set service-flow-monitor-
timer 1
```


Copyright© 2022 Ciena® Corporation

```
traffic-profiling dynamic-profile set microflow-monitor-timer 1

traffic-profiling dynamic-profile set greedy-microflow-
monitor-timer 1

traffic-profiling dynamic-profile set microflow-batch-monitor-
count 20

traffic-profiling dynamic-profile create name test

traffic-profiling dynamic-profile set profile test service-
flow-cir 200000

traffic-profiling dynamic-profile set profile test service-
flow-utilization-threshold 100000

traffic-profiling dynamic-profile set profile test service-
flow-over-threshold-count 3

traffic-profiling dynamic-profile set profile test service-
flow-under-threshold-count 3

traffic-profiling dynamic-profile set profile test greedy-
microflow-cir 10000

traffic-profiling dynamic-profile set profile test greedy-
microflow-drop-threshold 1000

traffic-profiling dynamic-profile set profile test greedy-
microflow-under-threshold-count 3

traffic-profiling dynamic-profile set profile test microflow-
over-threshold-count 3

traffic-profiling dynamic-profile set profile test microflow-
under-threshold-count 3

traffic-profiling dynamic-profile set profile test microflow-
monitor-count 10

traffic-profiling dynamic-profile add profile test port 1 vlan
10

traffic-profiling dynamic-profile enable profile test

traffic-profiling dynamic-profile enable
```

2 Use the following CLI commands to check the behavior:

```
traffic-profiling show

traffic-profiling dynamic-profile show

traffic-profiling dynamic-profile show detected-microflows
```


Copyright© 2022 Ciena® Corporation

```
traffic-profiling dynamic-profile show greedy-microflows

traffic-profiling dynamic-profile show profile test

traffic-profiling dynamic-profile show profile test detected-
microflows

traffic-profiling dynamic-profile show profile test greedy-
microflows
```

Example 2

If the user wants to enable dynamic microflow detection on port 2 also for VLAN 10 traffic then he or she can use the following two options:

```
traffic-profiling dynamic-profile remove profile test port 1
vlan 10

traffic-profiling dynamic-profile add profile test port 1-2
vlan 10
```

(This is considered as a single association, resource friendly).

Or

```
traffic-profiling dynamic-profile add profile test port 2 vlan
10
```

(This is treated as a new association and is not recommended.)

Example 3

To avoid overlapping, the following sequences or reverse sequences are not allowed:

```
traffic-profiling dynamic-profile add profile test port 1 vlan
10

traffic-profiling dynamic-profile add profile test port 1 ->
Overlap
```

Or

```
traffic-profiling dynamic-profile add profile test port 1-2

traffic-profiling dynamic-profile add profile test port 1 ->
Overlap
```

Or

```
traffic-profiling dynamic-profile add profile test port 1-2
vlan 10
```


Copyright© 2022 Ciena® Corporation

```
traffic-profiling dynamic-profile add profile test port 1 ->
Overlap
```

—end—

Copyright© 2022 Ciena® Corporation

Copyright© 2022 Ciena® Corporation

Multicast services configuration

Multicast services deliver data frames to multiple destinations by means of frame replications. Frame replications take place in a switch that is directly connected to a multicast server or indirectly connected to a multicast server through a router or switch. In multicast services, frame replication can be to selected members of the flood domain, decided either statically through channel stream or dynamically using Internet Group Management Protocol (IGMP) snooping. This allows network devices to discern multicast traffic and act on it according to the local content rules and processes for the platform.

Multicast services in the network comprise these components:

- one or more multicast servers
- one or more hosts
- routers that replicate the multicast packets at L3 and send to multiple interfaces
- switches at L2 that forward multicast frames to multiple interfaces towards the hosts, that is, the multicast termination points

Copyright© 2022 Ciena® Corporation

This figure shows multicast services in the network.

Figure 9-1
Multicast services in the network

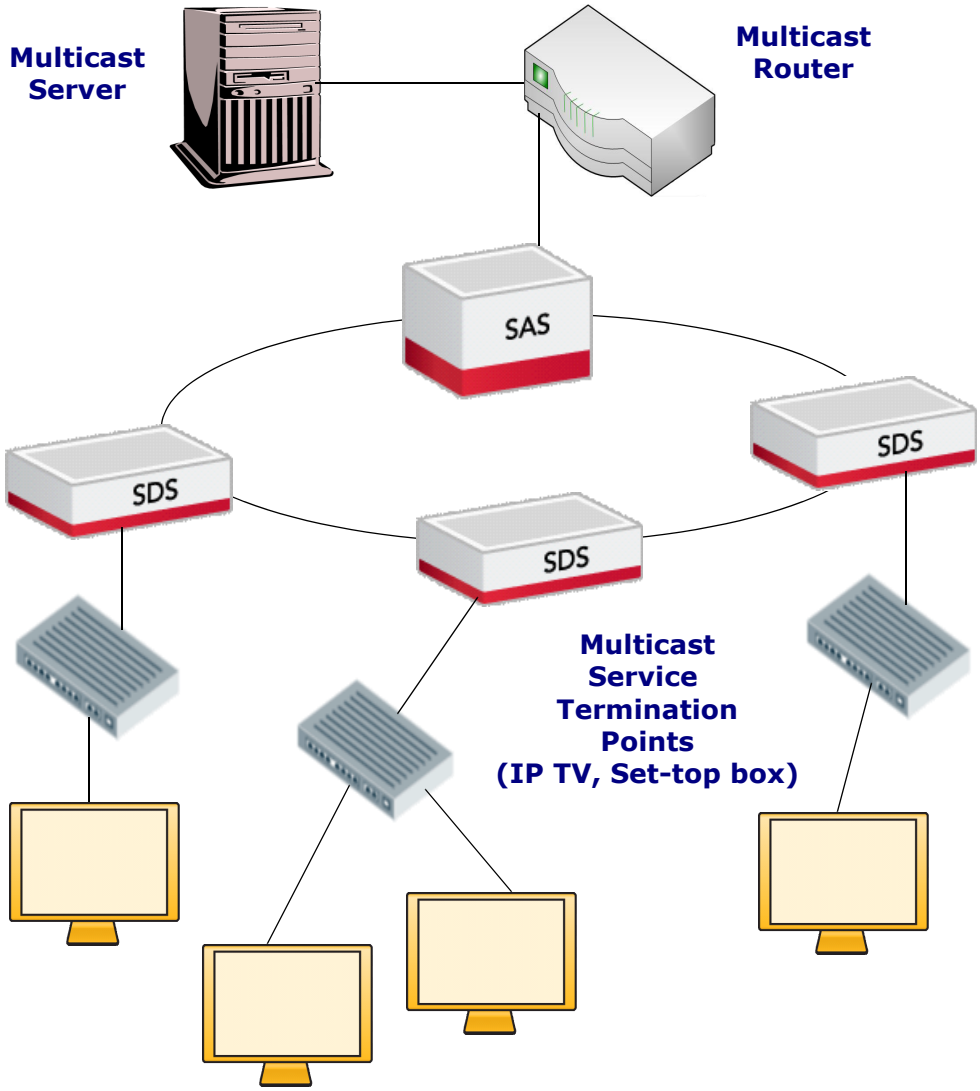

Routers determine the multicast streams to send to a network interface by means of IGMP. IGMP is used by IP hosts and multicast routers to establish and maintain membership in multicast groups.

Copyright© 2022 Ciena® Corporation

This table lists IGMP messages.

Table 9-1
IGMP messages

Message	Description
Report/Join	Indicates that a host wants to receive a specific multicast stream.
Leave	Indicates that a host wants to leave a specific multicast stream.
Query	Indicates that an IGMP querier is determining which hosts are on the interface. An IGMP querier is typically the router.

The switch establishes the multicast group filters either statically or dynamically.

Periodic queries are sent at a time interval. A group is removed after a timeout period if no join is heard back from the query.

Topics for multicast services configuration are:

- "IGMP snooping" on page 9-3
- "Enhanced features" on page 9-4
- "Multicast operations" on page 9-5
- "Statistics" on page 9-12
- "IGMPv3" on page 9-14
- "Services supported" on page 9-21
- "IGMP interworking" on page 9-22
- "Multicast-services attributes" on page 9-31

IGMP snooping

The switch maintains a list of multicast groups and a list of member interfaces for each group. Interfaces are added and removed from the multicast group based on IGMP join and leave messages received from a port.

The switch periodically updates the multicast group membership list. Interfaces are removed from the multicast group if a join message is not received from the interface after a timeout period.

If a multicast packet is received and the destination address is found in the list of multicast groups, then this packet only egresses the Interfaces that are members of the group. All other multicast packets flood to all Interfaces.

Copyright© 2022 Ciena® Corporation

The switch identifies one Interface as the router port. The selection of the router Interface is determined by general queries snooped by the switch.

IGMP join and leave messages are filtered by the switch. The switch forwards the first join message and the last leave message.

Enhanced features

This table lists the enhanced features of multicast services that are supported by 39XX/51XX switches.

Table 9-2
Enhanced features for multicast services

Enhanced feature	Description
Proxy reporting	The switch replies directly to the router in response to an IGMP query message. When a query is received, the switch sends IGMP membership reports (for multicast groups that have at least one downstream host) to the router based on the list of active multicast groups. The reports are sent to the router at a constant rate, which reduces the bursting that otherwise occurs. This feature is always on when IGMP snooping is enabled.
Proxy query	The switch sends the query message to one interface at a time. The delay between interfaces can be configured by the network operator. The delay balances the load on the CPU which reduces the channel-change latency. This feature is always on when IGMP snooping is enabled.
Query engine	The network operator configures the switch to generate IGMP query packets at regular intervals just like a router. Multicast servers can connect directly to the multicast network without going through routers.
Inquisitive leave	This feature supports network configurations that can have multiple IGMP hosts on a single interface. Two or more hosts are "tuned" into a single multicast group at the same time on the same interface. If the switch responds immediately to an IGMP leave message for this group, service to the other host or hosts tuned into the group is interrupted. The switch sends one group-specific query to the interface with max-response set to last-member-query-interval, and a timer is created with the same timeout value. As soon as an IGMP membership report is received on the port for the specific group, the timer is canceled. The interface is removed from the group if the timer expires. This feature can be enabled on any IGMP snooping-enabled L2 forwarding domain.

Copyright© 2022 Ciena® Corporation

Table 9-2
Enhanced features for multicast services

Enhanced feature	Description
Channel stream	A channel stream is a collection of user-defined multicast groups. The switch always subscribes to these groups, even if there are no hosts joining these groups. This feature can be used to reduce the latency experienced by the end user by forwarding some multicast streams far down into the network before they are requested by any IGMP host.
Unknown multicast filtering (UMF)	This feature is used to filter multicast packets with unknown destination addresses. The destination address is considered unknown if it cannot be found in the list of active multicast groups. UMF can be enabled on any multicast-enabled VLAN and virtual switch. The system tracks the setting for each L2 forwarding domain and passes this information down to the forward engine.
IGMP forking	IGMP forking is the process of forwarding, that is, forking, the individual raw join and leave messages that ingress the UNI port to the associated subscriber or data service, that is, VLAN, egressing the NNI ports. These raw subscriber-associated IGMPv3 and IGMPv2 PDUs provide opportunities for northbound systems to perform additional statistics tracking and QoS treatments. IGMP snooping must be enabled for forking to work. IGMP forking is not supported on virtual switches. See "Configuring IGMP forking with VLAN translation" on page 9-38.

Multicast operations

Multicast operations comprise:

- "Multicast forwarding domains" on page 9-6
- "Multicast interface" on page 9-6
- "Multicast traffic filters" on page 9-6
- "Multicast servers and routers" on page 9-7
- "Server topology" on page 9-8
- "IGMP query engine" on page 9-9
- "Router IP address range" on page 9-9
- "Channel stream" on page 9-10
- "Rapid recovery mode" on page 9-11
- "Layer 2 layer 3 mismatch behavior" on page 9-11
- "IGMP groups" on page 9-12

Copyright© 2022 Ciena® Corporation

Multicast forwarding domains

39XX/51XX switches use VLANs or virtual switches as the L2 forwarding domain. The characteristics of the VLAN or virtual switch are combined with the multicast group IP address to create unique entries in forwarding tables.

Multicast interface

SAOS uses ports and/or MPLS-VCs as multicast interfaces. When the network operator enables multicast-services on a forwarding domain, all interfaces in the forwarding domain are treated equally, that is:

- multicast streams can ingress any interface and egress on any other interfaces
- all multicast groups in the channel stream egress all interfaces
- traffic for multicast groups in the channel stream egresses all interfaces except the one the traffic is received on
- any interface can become the router interface

The network operator can configure two sets of interfaces:

- server interface
- channel stream interface

Multicast traffic filters

Multicast features, for example, channel stream and IGMP snooping, are methods for creating and deleting multicast traffic filters.

The Example of a multicast traffic filter table provides an example of a multicast traffic filter table. In this example, the L2 forwarding domain has a total of eight interfaces which are numbered 1 to 8. There are two VLANs on the switch: VLAN-1 and VLAN-2. Interfaces can belong to more than one VLAN. Also note that each entry in the table is uniquely identified by a key which comprises the VLAN and multicast destination IP address.

Copyright© 2022 Ciena® Corporation

Note: Multicast addresses are specified as Layer 3 IP addresses.

Table 9-3
Example of a multicast traffic filter table

Interface	VLAN (key)	Destination address (key)	Egress interface list	Notes
1	1	225.0.2.1	1,2,3,4	This is the typical multicast group entry.
2	2	225.0.2.1	5,7	This entry has the same destination address as entry 1 but has a different VID and egress interface list. This filter only affects the traffic on VID 2.
3	1	225.0.2.2	6	This entry is a multicast group with a single egress interface.
4	1	225.0.2.3	none	This entry is a multicast group with no member interfaces. This group can be in the linger state.
5	1	225.0.2.1	1,2,3,4,6,8	This entry is a channel stream group. Interface 5 is excluded from channel stream.
6	1	225.0.2.2	1,2,3,4,6,8	This entry is another channel stream group. All channel-stream groups have the same egress interface list.
8	1	Range: 225.0.2.0 to 225.0.2.255	none	This entry is the UMF filter. The destination address is the complete range of multicast addresses. This entry is the last entry in the table. It is the default entry when no other match is found in the table.

Multicast servers and routers

A multicast server is considered to be any network device that forwards multicast packets in the network. A multicast router is a special type of multicast server. Routers are generally L3 devices running IGMP and generating IGMP query messages. Any number of multicast servers can be used on a forwarding domain at the same time, but only one device can act as the IGMP query engine at any given time.

Copyright© 2022 Ciena® Corporation

Server topology

SAOS supports Distributed or Centralized Server topology. This table lists multicast delivery models for multicast services delivered over L2 networks.

Table 9-4
Multicast delivery models for multicast services delivered over L2 networks

Multicast delivery mode	Description
One-to-many	One server sends multicast streams to many hosts on the network. Membership reports are sent to a single server interface. Example: IPTV. This is a centralized delivery model.
Many-to-many	Many servers send multicast streams to many hosts. Membership reports are sent to all server interfaces in the network. Example: Video conference. This is a distributed delivery model.
Many-to-one	Many servers send multicast stream to one host. Membership reports are sent to all server interfaces in the network. Example: surveillance. This is a special case of many-to-many.

In the centralized server topology, all multicast services are delivered from one L2 device to the rest of the network. Any number of multicast servers can be used, but all multicast streams must flow downstream from one central point and all membership reports flow upstream to this central point.

In the distributed server topology any number of multicast servers on the network can interface with any device on the network.

This table describes IGMP behavior for server topologies.

Table 9-5
IGMP behavior for server topology

IGMP function	Centralized	Distributed
Receive General Query	Reply to the router. Send membership reports for all active multicast groups.	Reply to the router. Only send membership reports for channel stream groups.
Send Report (join/leave)	Send to the router interface only.	Send to all server and router interfaces.

Copyright© 2022 Ciena® Corporation

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 needed in a converged service network. To eliminate the need for a separate IGMP router and speed up multicast response times, the network operator can configure service delivery switches to use a built-in IGMP query engine.

When the query engine is disabled, the switch does not send any query message unless a router interface is learned. The router interface is learned by snooping periodic IGMP query messages sent by the router. Once the router interface is learned, the switch sends query messages until a timeout occurs waiting for the next query from the router. The switch saves the source IP address and source MAC address of the router and uses these addresses whenever the switch needs to send a query message to the network.

When the query engine is enabled, the switch uses the IP source address configured by the user whenever it needs to send a query message to the network. The switch starts sending general queries as soon as query-engine and igmp-snooping are both enabled and they continue until the configuration is changed. In this mode, the switch does not respond to a query received from another device on the network unless that device has a lower IP address. If the system receives another query with a lower source IP address, the query engine goes inactive. The switch then learns the router interface as previously described.

If you enable query-engine on two or more devices on the same network, both devices send queries to the network but only one device, that is, the one with the lower IP address, receives replies.

> ***Note:*** When the server topology is set to centralized, only one switch on the network can act as the master querier and all multicast servers must interface directly with the master querier.

Router IP address range

The network operator can define a single range of multicast addresses for the router. If the network operator defines this address range, the switch does not send a membership report to the router for any multicast group that does not fall inside this address range. If the network operator does not define this address range, the switch forwards all membership reports to the router regardless of the group address.

When two or more routers are used on the same network, that is, the server-topology attribute is set to distributed, the router address range must remain undefined so that all join messages are forwarded to all routers.

Copyright© 2022 Ciena® Corporation

Channel stream

Any network running IGMP is bound to experience a certain amount of latency between the IGMP join message and the arrival of the multicast group itself. Channel stream reduces this latency and reduces the processing load on the router. Channel stream defines a set of multicast groups that are flooded to all interfaces in the forwarding domain, which brings multicast services far down into the network, that is, closer to the subscriber, before any host joins the groups. Latency is reduced because the IGMP join message does not need to propagate all the way to the router. If channel stream is configured such that multicast groups flow all the way to the switch, then the switch is the only device that must process the join message to forward the multicast group to the subscriber.

The channel stream must be configured before any multicast groups flow. When the multicast group is added you see a channel stream group. Multicast groups can be added or removed from this channel stream at any time. As multicast groups are added to the channel stream, they immediately begin to flow to all interfaces in the forwarding domain.

The network operator can define a list of channel-stream interfaces. By default, all interfaces in the forwarding domain belong to the channel stream interface list. The network operator can remove any interface from this list.

Channel stream is used with UMF. If a multicast stream arrives at the switch before the multicast group is added to channel stream, the switch floods or drops the stream depending on the setting for UMF. After the group is added to channel stream, the group floods to all channel stream interfaces regardless of the UMF setting. UMF must be set to drop on the channel stream device and all downstream devices. This behavior occurs when IGMP snooping is enabled.

Channel stream can be used with or without IGMP snooping. In most cases, snooping must be enabled when channel stream is used, but it is not required to do so. If multicast services originate from a router, then IGMP snooping must be enabled in order for the switch to send IGMP membership reports to the router. Channel stream can be used in conjunction with UMF as a way to selectively forward a subset of the multicast groups from a multicast server to the rest of the network without IGMP snooping.

If a multicast group is added to channel stream and the source of the multicast stream is the router, then the switch becomes a member of the multicast group, and it behaves just like any other group member. The switch immediately sends an IGMP join message to the router for this new multicast group, and it responds to IGMP queries as long as the multicast group is a

Copyright© 2022 Ciena® Corporation

member of channel stream. The switch forwards the multicast group to all channel stream interfaces. Any IGMP membership reports received by the switch from downstream devices are dropped.

The switch does not disrupt service to downstream devices when a multicast group is added or removed from channel stream. If a multicast group is already active on the switch when the group is added to channel stream, the switch adds more interfaces to the multicast group. Similarly, when a multicast group is removed from channel stream, the switch continues to flow the multicast group to all interfaces that subscribe to the group. The switch does not keep track of group membership while the group is part of channel stream. When the group is removed from channel stream, the switch sends a group-specific query to all interfaces in the network to determine the group membership. If no hosts respond to the query in the timeout, the group goes into the linger state.

It is not necessary to define the source interface for channel stream. It is also not necessary for all groups in the channel stream to have the same source interface. Some streams can ingress from a router while others ingress directly from a video server. The switch sends join messages to the router for all multicast groups in the channel stream.

Rapid recovery mode

Rapid recovery mode is a configuration option which determines how general queries with a source IP address of 0.0.0.0 are processed or transmitted. When this setting is enabled, the switch can process and transmit general queries with a 0 IP source address. If the setting is off, general query packets are not transmitted or processed.

This setting is useful when the switch loses its router interface. Rapid recovery enables the switch to transmit general queries with an IP source address of 0 on all the layer 2 forwarding domain (L2FD) interfaces. New router interfaces are learned faster without having to rely on configured query interval timers.

When rapid recovery mode is enabled and the device receives a general query with an IP source address of 0, and the general query is received on the router interface, the router interface is marked as unknown. General queries are sent out to all neighbors. If the device knows its router interface, then the router replies with a general query and provides the details of the active router. For anything else, the general query is forwarded to all the neighbors.

Layer 2 layer 3 mismatch behavior

If a switch receives a general query (IGMP group address of 0.0.0.0) then a check is made to ensure that the DA MAC is 01005e000001. If not, the message is a L2/L3 mismatch message and the necessary event is generated. Join and leave messages are not validated.

Copyright© 2022 Ciena® Corporation

IGMP groups

SAOS switches support up to 1024 multicast groups, as shown in this table.

Table 9-6
IGMP groups

Platform	Number of IGMP groups
3903, 3903x, 3905, 3906	253
3942	1022
5142, 5160	1021
3924, 3926, 3928	250

Note: Multicast and ENNI virtual switches use the same hardware resources. As a result, the maximum amounts specified may not be achieved if there is an ENNI virtual switch on the same device. In this case an error is generated.

Statistics

The switch keeps IGMP statistics. This table lists the statistics collected for each forwarding domain.

Table 9-7
Statistics collected for each forwarding domain

Statistic	Description
IGMPv2 Membership Reports	Number of IGMPv2 reports received by the system.
IGMPv3 Membership Reports	IGMPv3 membership report Rx also tracks by individual report type: • IS_IN Reports Rx • IS_EX Reports Rx • TO_IN Reports Rx • TO_EX Reports Rx • Allow Reports Rx • Block Reports Rx
Leave Reports	Number of leave messages received, with and without errors.
Queries Received	Number of query messages received, with and without errors.

Copyright© 2022 Ciena® Corporation

Table 9-7
Statistics collected for each forwarding domain

Statistic	Description
Query Discard	Number of query messages discarded: query packet received on non-server interface; query packet received before the reply to the last query is finished.
Router Timeouts	Number of times the router times out without receiving a query.
Query Reply Timeouts	Number of timeouts that occurred while sending membership reports to the router (error). This is incremented when there is no reply to a query (query-reply sequence) before the router times out (min(minQueryResponseInterval, maxResponseTime)).
Router Discards	Number of invalid IGMP frames from the router that have been discarded, for example, join message.
Host Discards	Number of join and leave messages from hosts that have been discarded, including: • snooping disabled • invalid multicast group address • ingress interface was the router interface (centralized mode only) • group is not active (leave packet only) • ingress interface does not belong to the forwarding domain • state of the ingress interface is not forwarding
Checksum Errors	Number of IGMP PDU checksum errors.
L2/L3 Mismatch	Number of mismatches of L2/L3 mismatch. This is when the L2 destination address does not match the destination multicast IP address.
Unknown Packet Type	Number of IGMP packets received that are of type unknown.
Exceeded Resources	Number of times a join message is received but the system cannot allocate a new group due to resource limitation.
Query Discarded because rapid recovery mode-off	Number of queries discarded because rapid recovery mode is off.
Compatibility Mode Discards	The IGMP version did not match the configured settings.

Copyright© 2022 Ciena® Corporation

IGMPv3

IGMPv3 adds support for source filtering for a system to report interest in receiving packets from

- only specific source addresses
- all but specific source addresses that are sent to a particular multicast address

This information may be used by multicast routing protocols to avoid delivering multicast packets from specific sources to networks with no interested receivers, or to select specific sources for a multicast stream when the multicast address is not unique between the streams.

When IGMP snooping is enabled on a switch, supporting IGMP v3 means that source filtering multicast is supported. The switch forwards all received membership reports on the interfaces when an “include source” and “exclude source” membership report is received.

See these sections related to IGMPv3:

- “IGMPv3 messages” on page 9-14
- “IGMP snooping” on page 9-15
- “Membership reporting” on page 9-15
- “Interface state” on page 9-16
- “No host suppression” on page 9-18
- “Queries” on page 9-18

IGMPv3 messages

To support source filtering multicast, the source list and associated attributes must be added into some of the IGMP messages. IGMPv3 modifies the IGMP query message and the IGMP membership report message to accommodate this.

IGMPv3 query contains the general query and group-specific query as does IGMPv2. IGMPv3 adds a group-and-source-specific query message which includes the number of sources and the list of source addresses. If IGMP snooping is enabled on a virtual switch or VLAN, the proxy querier must support the new message format.

The IGMPv3 membership report is a new message 0x22 type that contains these attributes in addition to the IGMPv2 report:

- Multicast DA address
- Number of group records (M)
- List of group records with each containing a record type, number of sources and list of sources. The record type contains one of:

Copyright© 2022 Ciena® Corporation

  - Current-State: Mode-Is-Include (IS_IN) or Mode-Is-Exclude (IS_EX)
  - Filter-Mode-Change: Change-To-Include-Mode (TO_IN) or Change-To-Exclude-Mode (TO_EX)
  - Source-List-Change: Allow-New-Source (ALLOW) or Block-New-Source (BLOCK)
- Number of sources (N)
- List of sources (source address [1] to source address [N])

When IGMP snooping is enabled on a switch, proxy reporting must support the message format.

IGMP snooping

For IGMPv3, the addition of source address records means that the multicast group entry must remember the source address records for each host to ensure that they are aggregated together and correctly forward the Join northbound to the multicast router and respond to queries. Each group entry record also contains a copy of the source address records included in the Join and the mode associated with those records.

Membership reporting

Group membership reporting aggregates together the source address lists and their modes to generate a list that matches the requirements of the individual hosts subscribed to the multicast group. Source address lists may specify either an Include or Exclude mode. Include mode is used when the host wants to subscribe to a multicast stream from specific sources so that it includes the list of source address that it accepts the stream from. The host may also wish to subscribe to a multicast stream except from specific sources so it Excludes those sources from the source by means of the source address list. When the Join or Leave northbound are forwarded, the various hosts that have subscribed to the group are aggregated together to resolve conflicts between their source address lists and the modes of those lists.

When IGMPv3 membership reports are processed, individual group records of the report are processed. Any report-up-stream per group record is also sent. IGMPv3 membership reports sent to the router have only one group record.

> ***Note:*** Conflicts can occur during the aggregation of the lists. These conflicts are resolved by including a stream, if needed. It is better that one host ignores a stream from a source it wanted to Exclude rather than fail to deliver a stream to a host that wanted to Include it.

Copyright© 2022 Ciena® Corporation

Interface state

This table describes interface state attributes.

Table 9-8
Interface state

Attribute	Description
Group address	Multicast address for this group entry
Filter mode	Set to Include or Exclude.
Source list	A set of source addresses, such as {a, b, c, d, e} where a, b, c, d and e are each a source IP address.

For IGMPv3, an interface can be in reception state or not in reception state, as in IGMPv2. Reception of a query and a change of the interface reception state triggers IGMPv3 protocol actions on an interface. When the behavior is similar to IGMPv2, IGMPv3 specifies a new report format to handle Source Specific Multicast (SSM) information. In IGMPv3, when an interface is in reception state of a multicast group, it contains more attributes: filter-mode (include or exclude) and source list. A change in reception state can be a filter mode change, a source list change or both (covered by the filter mode change).

When receiving a group-and-source-specify query, the proxy reporting builds a membership report on a per-group interface list of sources. An interface of a multicast group with state INCLUDE (X) means that frames destined for the multicast group from any source addresses specified in X are forwarded to the interface. Similarly, an interface of multicast group with state EXCLUDE (X) for group G means to block the frames destined to G from any source addresses specified in X. The current state (IS_IN, IS_EX) group record is contained in the membership reports in response to queries.

When an IGMP snooping enabled virtual switch receives an interface reception state change, such as receiving a membership report from an interface that contains a state change record, the group filter must be modified and a state change report must be sent to the router.

Copyright© 2022 Ciena® Corporation

This table shows how the new state-change record is computed. the mapping.

Table 9-9
State change record sent

Old state	New state	State-change record sent
INCLUDE (A)	INCLUDE (B)	ALLOW (B-A, BLOCK (A-B)
EXCLUDE (A)	EXCLUDE (B)	ALLOW (A-B), BLOCK (B-A)
INCLUDE (A)	EXCLUDE (B)	TO_EX (B)
EXCLUDE (A)	INCLUDE (B)	TO_IN (B)

If the computed source list for an ALLOW or a BLOCK state-change record is empty, that record is not included in the report message. The state-change report is retransmitted one or more times at chosen intervals to ensure that the report is not missed by one or more multicast routers.

The IGMPv3 proxy querying/reporting figure shows an example of IGMPv3 proxy querying/reporting.

Note: The gray circles in this figure represent ports.

Figure 9-2
IGMPv3 proxy querying/reporting

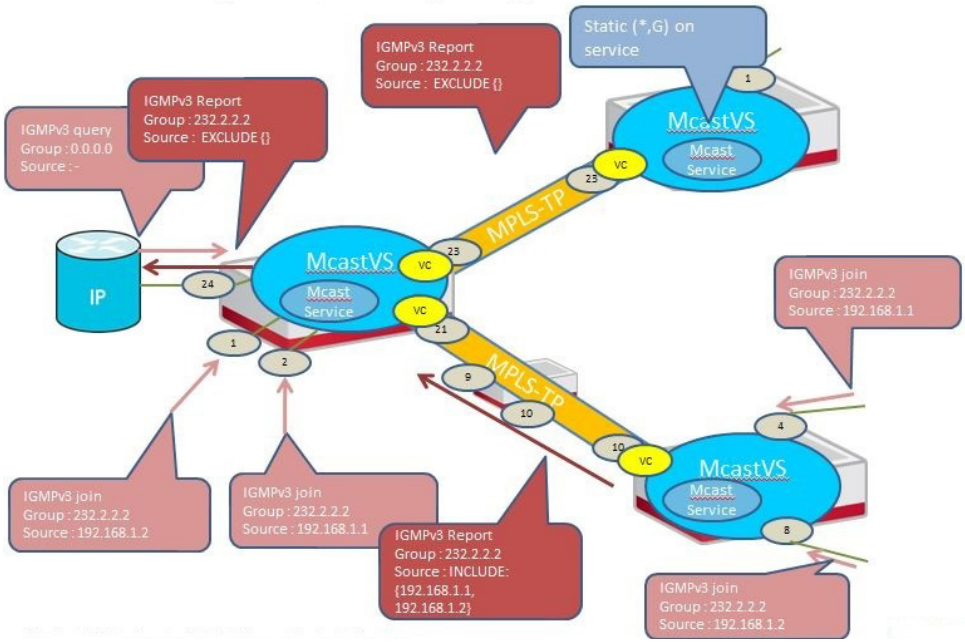

Copyright© 2022 Ciena® Corporation

No host suppression

Proxy query and proxy reporting remove the suppression requirement. One reason for this is that some routers want to keep track per host status for billing purposes. When receiving interface state change messages and updating the local database and multicast filter configuration, the message must be forwarded to the router. When forwarding host messages, the correct version of the message format must be used.

Queries

Sending a general query and a group-specific query is the same in IGMPv3 as it is in IGMPv2. For IGMPv3, the proxy query must send a group-and-source-specific query to the logical interface when it receives a change of interface state (filter mode or source list change) indicating the member on that interface is leaving the specific group and specific source.

Source records, filter mode and timers

This table describes how the system maintains source records to keep up-to-date membership reporting.

Table 9-10
Source records, filter mode and timers

Source records, filter mode and timers	Description
Group Filter Mode	Each multicast group per interface has a filter mode. The group filter mode is per multicast group. Providing that one interface of the multicast group is in EXLCUDE filter mode, the group filter mode is EXCLUDE.
Source Record	Source record lists the sources per multicast group. The source record's meaning depends on the group filter mode. For example, EXCLUDE {} is an empty source record and indicates that none of the sources are blocked.
Source Timer	A source timer is a timer per source of a source record. When the timer is running, there is at least one member from an interface that wants to receive multicast frames from this source. The source timers are updated when receiving a report containing the source in a group record. When the timer expires, no members want to receive the frames from that source.
Group Timer	The group timer is only used when a group is in EXCLUDE mode. It represents the time for the filter-mode of the group to expire and switch to INCLUDE mode. A group timer is a decrementing timer with a lower bound of zero kept per group per attached network. Group times are updated according to the types of records received.

Copyright© 2022 Ciena® Corporation

This table summarizes the IGMPv3 source specific forwarding rules.

Table 9-11
Source specific forwarding rules

Group filter-mode	Source timer	Action
INCLUDE	Timer is greater than 0.	Traffic is forwarded from the source.
INCLUDE	Timer equals 0.	Traffic stops forwarding from the source and the source record is removed. If there are no more source records for the group, the group record is deleted.
INCLUDE	Empty source list	Traffic is not forwarded.
EXCLUDE	Timer is greater than 0.	Traffic is forwarded from the source.
EXCLUDE	Timer is equal to 0.	Traffic is not forwarded from the source. The record is not removed for converting it back to INCLUDE.
EXCLUDE	Empty source list.	Traffic is forwarded from the source.

This table summarizes the IGMPv3 group timer rules.

Table 9-12
Group timer rules

Group filter-mode	Group timer	Action
EXCLUDE	Timer is greater than 0.	At least one member is in EXCLUDE mode.
EXCLUDE	Timer equals 0.	There are no more listeners to the group. If all source timers have expired, the group record is deleted. Otherwise, switch to IN filter mode using those sources with running timers.

Receiving current-state record

As current-state records are in response to a general query, the source list is updated and the forwarding engine is configured. The group filter mode is changed in special cases. These actions are summarized in RFC 3376. The router state is the group state kept in the multicast group.

Receiving filter-mode change or source-list-change record

The actions to receive a filter-mode change or a source-list change record are summarized in RFC 3376.

Copyright© 2022 Ciena® Corporation

Switching router filter mode

When a group filter mode is EXCLUDE and the group timer expires, the group filter mode transfers from EXCLUDE to INCLUDE, and the source list is updated. For example, if the router state of a group is EXCLUDE (X, Y) and the group timer expires, its router state is changed to INCLUDE.

IGMP version interoperability

IGMPv3 hosts and routers interoperate with hosts and routers that have not been upgraded to IGMPv3. This interoperability is supported in the proxy reporting and proxy querier by maintaining the proxy reporting and proxy query parameters.

Proxy reporting receives queries, responds to queries, and sends a report to routers/queriers when there is a change. The proxy report keeps a host querier compatibility mode for each interface where there are queries received, and an associated timer. By default, the host querier compatibility mode is IGMPv3. When a query received from an interface is of a lower version, the host compatbility mode changes to IGMPv2, and an older version querier present timer runs. The IGMPv2 timer resets for every query of version 2 received. If the older version querier present timer expires, the host querier compatibility mode is automatically reverted to IGMPv3.

When an interface's host querier compatibility mode is IGMPv2, the interface has to run the IGMPv2 protocol. For example, reports sending out to that interface have to use the IGMPv2 format. When the interface's host querier compatibility mode is IGMPv3, it must run the IGMPv3 protocol.

The host querier compatibility mode is updated when the older version querier present timer expires or when it receives queries with a lower version. When a host changes its querier compatibility mode, all pending response and retransmission timers are canceled.

The proxy query must keep a group querier compatibility mode for each multicast group, and an associated older version host present timer. By default, the group querier compatibility mode is IGMPv3. When a group receives a lower version report, the group querier compatibility mode changes to IGMPv2, and the older version host present timer starts. For every lower version report received, the timer is reset. If the timer expires, the group querier compatibility mode automatically reverts to IGMPv3. A timer expires when lower versions are not received for a specified time period. When the group querier compatibility mode is IGMPv2 or IGMPv3, queries sent to all interfaces of this group use the IGMPv2 and IGMPv3 protocol respectively.

Copyright© 2022 Ciena® Corporation

Different version hosts can be attached on a port. Version hosts can participate in a different version of host reports per different multicast members. This results in different members being in different querier compatibility modes. Both versions of the general query are used.

When the group querier compatibility mode is IGMPv3, a router acts using the IGMPv3 protocol for that group. When the group querier compatibility mode is IGMPv2, a router internally translates IGMP messages for that group to their IGMPv3 equivalents.

Services supported

This table outlines the services multicast supports.

Table 9-13
Services supported

Service type	IGMP supported	SSM supported	Interface types supported
QinQ VS	Yes	Yes	Logical ports
VLAN	Yes	Yes	Logical ports
VPLS/VPWS VS	Yes	No	Logical ports plus MPLS VC
VS with PBT-service	No	Not applicable	Not applicable
VS with sub-ports	No	Not applicable	Not applicable

Limitations

Multicast services over VS with CVID bundling are not supported. The recommended configuration for IGMP enabled VS is as follows:

- The VS must only have EVPL without CVID bundling.
- All UNI ports must have ingress POP and egress PUSH transforms.
- The IGMP module must always send packets towards the router as single S-TAG packets for QinQ.

Copyright© 2022 Ciena® Corporation

This table outlines the recommended actions if IGMP enabled VS is not configured correctly.

Table 9-14
Misconfiguration scenarios

Scenario	Action
CVID bundling is configured.	• Ingress action is still POP. • On egress, the last egress transform configured is used. • Single leave from one of the DSLAMs connected to a UNI port trigger removes the port from the FWD bitmap.
Transforms are not configured.	• IGMP module still sends single S-TAG tagged packets towards the router. • Double-tagged IGMP packets (MPLS payload) received on an MPLS VS NNI are not trapped to CPU.

Source specific multicast (SSM) is not supported for MPLS-based multicast traffic forwarding. For MPLS-VS-based multicast services with mesh-vc, IGMP treats all mesh virtual circuits as a single logical interface. For example, if one of the mesh virtual circuits is a router interface, IGMP treats all other mesh virtual circuits as belonging to the same logical router interface. As a result, it may discard any membership reports ingressing any of the mesh virtual circuits. If a spoke virtual circuit is the router interface, IGMP filters the membership reports on the mesh virtual circuits individually, and not as a single logical interface.

Enabling multicast services on any VLAN/QinQ L2FD traps all IGMP PDUs (belonging to other VLAN/QinQ L2FD) to CPU on all platforms. These packets are then forwarded by the CPU.

IGMP interworking

Internet Standard Multicast (ISM) is an entity that responds to snooped IGMP PDUs and accordingly controls the flow of multicast traffic. This section describes how ISM reacts to external network agencies such as port, operational state change, RSTP and MSTP changes, G.8032 channel blocks and MC-LAG events.

See these sections for related information on IGMP interworking:

- “Interface state change” on page 9-23

Copyright© 2022 Ciena® Corporation

- "Port state change" on page 9-23
- "RSTP and MSTP events" on page 9-24
- "G.8032" on page 9-24
- "LACP" on page 9-28
- "MC-LAG" on page 9-28
- "MPLS" on page 9-30

Interface state change

This table describes the steps required when there is an interface state change.

Table 9-15
Interface state change

Interface state	Steps to take
The interface state changes to down/ blocked.	• If the interface is the router interface, mark the router interface as unknown. • If the interface is the default router port, stop sending reports. • Otherwise, remove this interface from all multicast groups associated with this service instance.
The interface state changes to up/ forwarding.	• If the interface is the default router interface and the router interface is unknown, activate the default router interface. • If the interface is not the default router interface, send GQ on this port. Use the current proxy-query-source-address. If the router port is not known, then the proxy-query-source-address is zero. Otherwise, it is the router source address. In either case, this packet informs the other device of the information known.

Port state change

This table describes how ISM reacts to port state changes.

Table 9-16
Port state change

Port state	Steps to take
PortEvent_StateChange	ISM reacts when the port state changes to up or down.

Copyright© 2022 Ciena® Corporation

RSTP and MSTP events

This table describes how ISM reacts to RSTP events.

Table 9-17
RSTP events

RSTP events	Steps to take
RstpEvent_PossibleTopologyChangeDetected	• The master querier sends out a GQ two seconds after a RSTP/MSTP topology change is detected.
RstpEvent_PortStateUpdatedInternal	• Set the RSTP state to block. • Set the RSTP state to forwarding.

This table describes how ISM reacts to MSTP events.

Table 9-18
MSTP events

MSTP events	Steps to take
MstpEvent_PossibleTopologyChangeDetected	• The master querier sends out a GQ two seconds after a RSTP/MSTP topology change is detected. • MSTP is set up on the virtual-circuit VLAN (SVLAN) and MSTP takes care of any loops on the NNI. If there was a ring setup on the NNI side, MSTP stops the loops on the SVLAN.
MstpEvent_PortStateUpdatedInternal	• A map is sent to the affected VLANs. • ISM goes over the serviceInstance database to determine if multicast services are going over any of the affected VLANS. ISM handles blocked/forwarding transitions.

G.8032

ISM works with G.8032 to provide a rapid recovery of the router interface for each IGMP-enabled device present on the ring. ISM stops processing and transmitting IGMP PDUs on an interface channel block. The fault on the affected link is detected and the channel block is moved to each end of the faulted link.

> ***Note:*** The successful delivery of multicast stream (to the hosts) is re-established within 1 sec after a G.8032 topology change when rapid recovery mode is enabled on all the nodes.

Copyright© 2022 Ciena® Corporation

This figure shows the G.8032 ring working normally with rapid recovery mode disabled by default.

Figure 9-3
G.8032 working normally

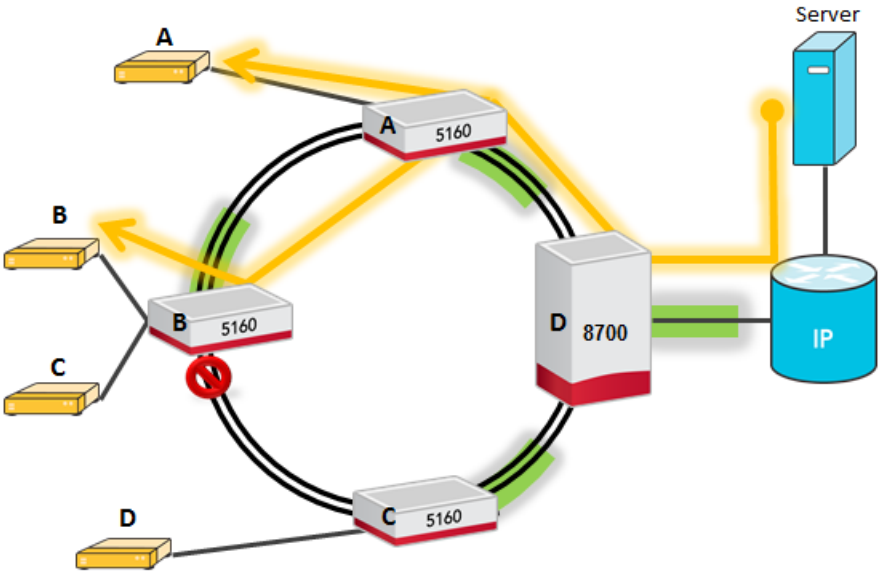

Copyright© 2022 Ciena® Corporation

A fault is detected and the channel block is moved to each end of the faulted link as shown in this figure.

Figure 9-4
Port state change detected

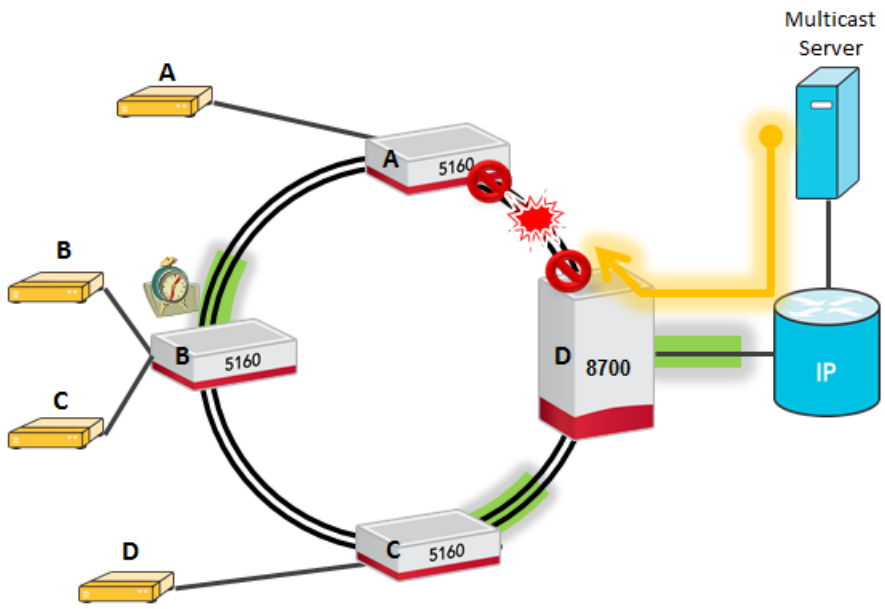

When the port state change is detected the router interface reference is removed from the faulty link. The router interface link BA for node B times out and is removed due to lack of queries from node A. Node C still has a valid link to the router and continues to forward the GQs through the CB link which is forwarded to the BA link.

Copyright© 2022 Ciena® Corporation

This figure shows that devices B and A re-learn the router interfaced based on the IGMP Query from device C.

Figure 9-5
Router interface removed

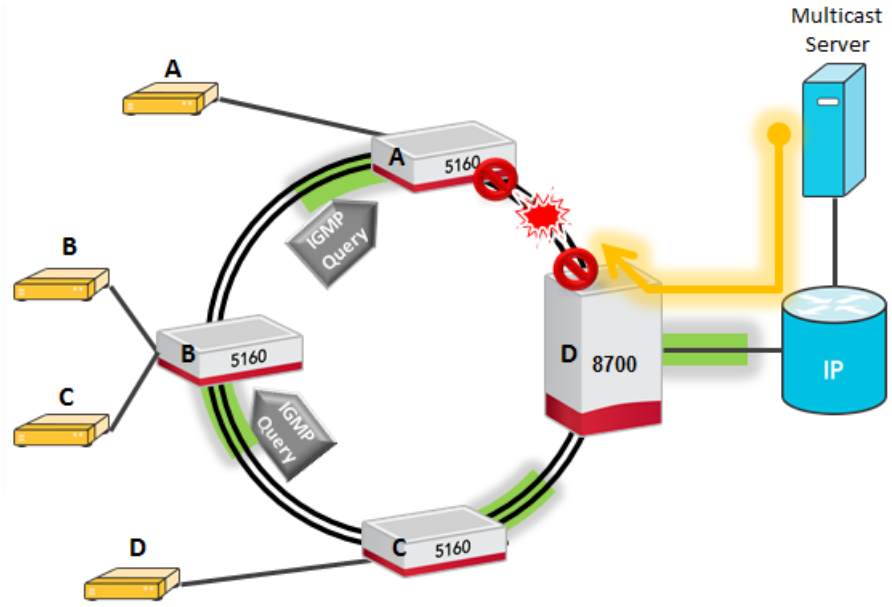

Copyright© 2022 Ciena® Corporation

Nodes B and A have learned the new paths to the router and traffic resumes along the new path, as shown in this figure.

Figure 9-6
New paths learned

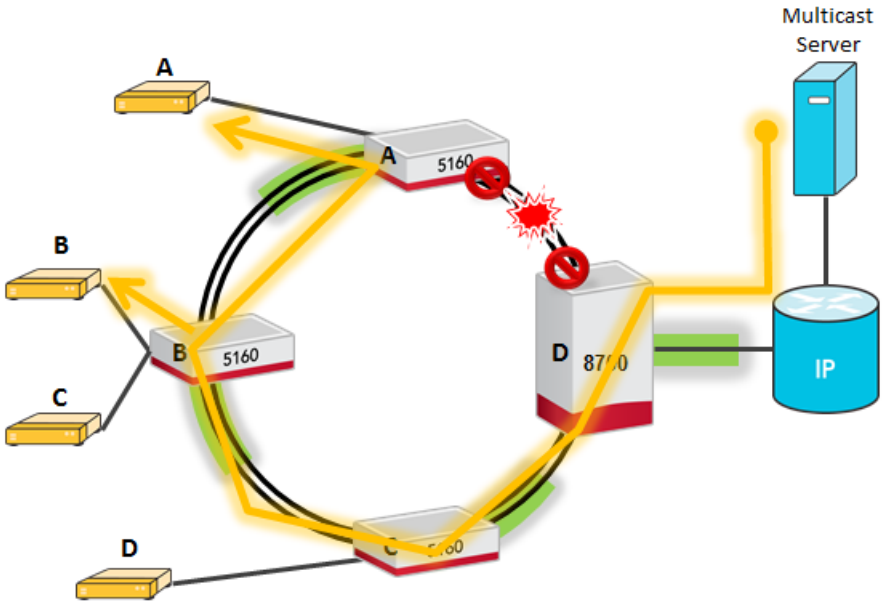

LACP

IGMP and multicast services are supported over aggregated ports. IGMP handles events such as port down and port up the same way it does for regular ports. To work correctly, IGMP acts on the notification of these events:

- Aggregation port is removed from the VLAN/virtual switch group
- Aggregation port deleted
- Aggregation port member ports have changed

These events change the port bitmaps of the relevant multicast groups. Once the appropriate bitmap is programmed, the hardware handles the load balancing and acts on member port failures.

MC-LAG

IGMP and multicast services are supported over MC-LAG. This figure shows MC-LAG working normally.

Copyright© 2022 Ciena® Corporation

Figure 9-7
MC-LAG working normally

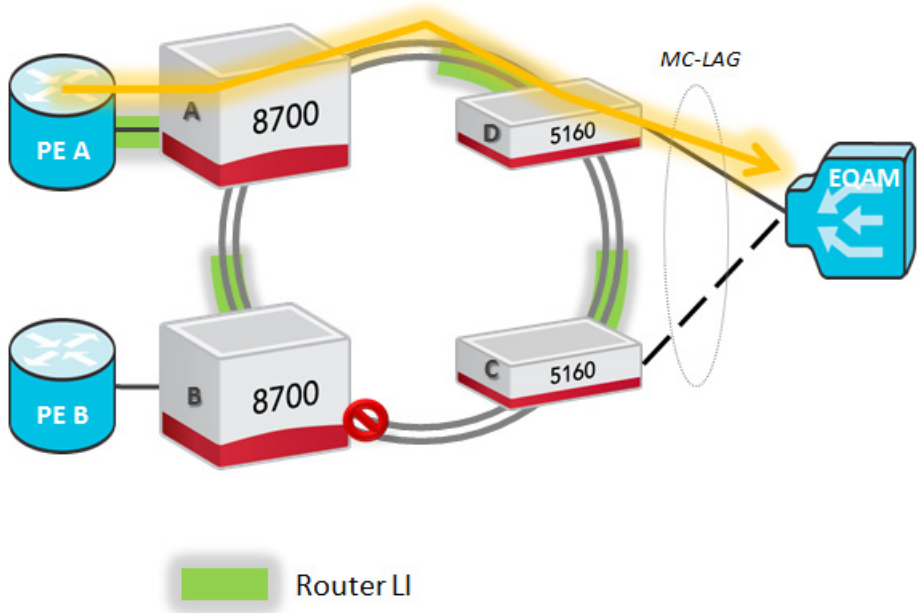

When the EQAM-facing logical interface (LI) of node D goes down, an MC-LAG switchover happens between nodes C and D, as shown in this figure.

Figure 9-8
MC-LAG switchover

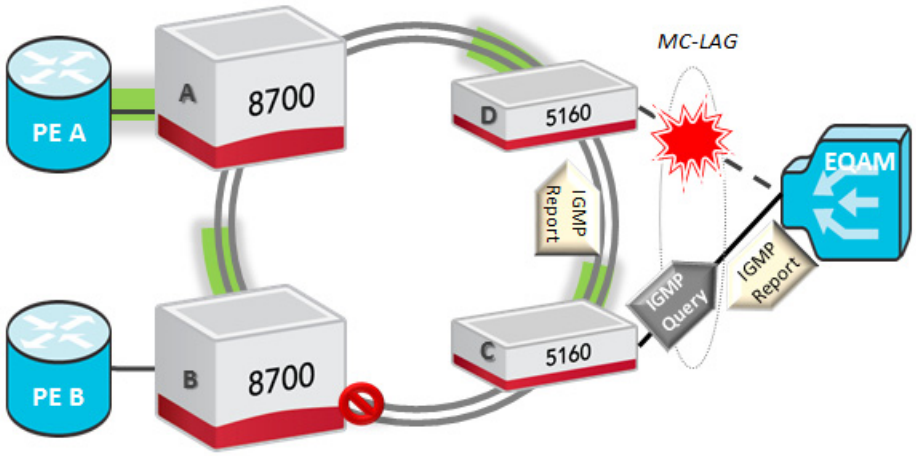

Copyright© 2022 Ciena® Corporation

For IGMP, this is an LI up event for the EQAM-facing LI on node C. A general query is sent to the LI followed by a host membership report from the EQAM-side. Node C then adds this LI to the multicast forwarding group. This LI of node C becomes the LI router for all downstream devices. Traffic flow then resumes through node C, as shown in this figure.

> ***Note:*** It is recommended that the customers using LAG with 3rd party vendors should switch to the standard mode as the 3rd party devices do not support the proprietary protection mode.

Figure 9-9
Traffic flow resumes

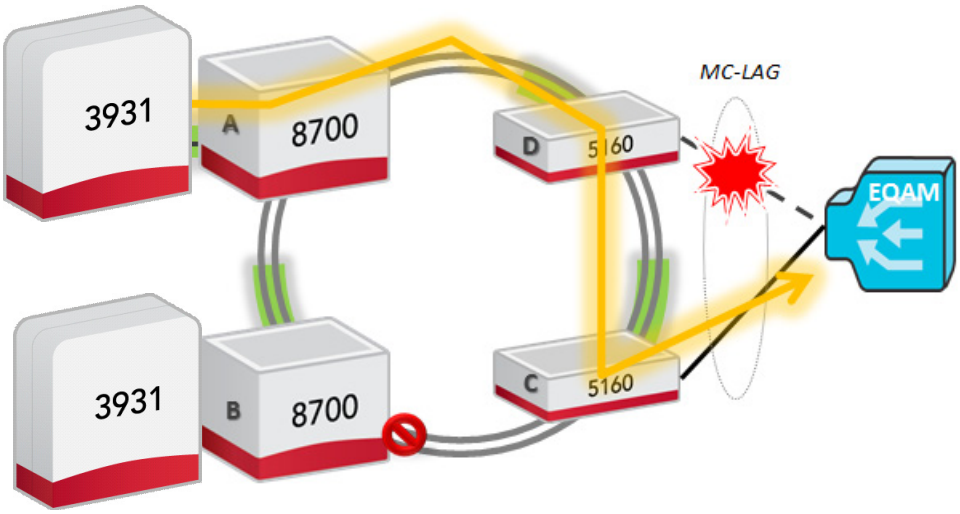

MPLS

The multicast service can be associated with a virtual switch, which can then be associated with virtual circuits.

There are two scenarios that can take place in MPLS.

- If the LI router is associated with a virtual circuit attachment that is being supported by working and protecting LSPs, and a LSP switchover occurs, recovery time depends on the LSP switchover time.
- If there are two pseudowire virtual circuits (active and standby), and if the LI router is associated with an active pseudowire fails, the multicast stream recovery time is the same as with G.8032. The LI router is removed from the active pseudowire virtual circuit and the rapid recovery process begins.

Copyright© 2022 Ciena® Corporation

- For multicast services over MPLS-VS with mesh mpls-vcs, IGMP considers all mesh VCs as a single logical interface. For example, if one of the mesh VCs is a router interface, then IGMP treats all the other mesh-VCs as belonging to the same logical router interface, and then discards any membership reports ingressing any of the mesh VCs. If a spoke VC is the router interface, then IGMP filters membership reports on the mesh-vcs individually and not as a single logical interface.

  Note 1: Multicast services over virtual switch with CVID bundling are not supported. The virtual switch must only have EVPLs without CVID bundling. All UNI ports must have ingress POP and egress PUSH transforms.

  Note 2: SSM is not supported over MPLS-based multicast traffic forwarding. SSM is not supported at a group level for virtual switch-based multicast services.

  Note 3: Sub-ports separately compete for the same internal resources as MPLS virtual circuits.

Multicast-services attributes

This table lists multicast-services attributes.

Table 9-19
Multicast-services attributes

Attribute	Default value	Description
active-linger-timeout <SECONDS: 0-300>	30	Specifies how long the group stays in the active state after all members leave the group. ***Note:*** This attribute is ignored in distributed mode.
linger-timeout	120	Starts when a leave message is sent to the router. The multicast filter is deleted when the timer expires.
querier-compatibility mode	v3	Specifies the IGMP version when Query Engine is on.
default-router-port <Port>	0	Specifies the default router port to use if the actual router port is not known. ***Note:*** This attribute is not used in distributed mode.
default-router-mpls-vc	0	Specifies the default router interface to use as a MPLS VC, if the actual router port is not known.
fork <on|off>	off	Indicates whether to fork IGMP frames.

Copyright© 2022 Ciena® Corporation

Table 9-19
Multicast-services attributes

Attribute	Default value	Description
last-member-query-interval <DECISECONDS: 10-100>	10	Specifies the maximum response time inserted into G-S queries sent in response to leave messages. Also the amount of time between G-S queries. ***Note:*** This attribute is Ignored when inquisitive-leave is disabled.
leave-mode <fast | inquisitive>	fast	Specifies how to handle leave messages.
linger-timeout <SECONDS: 10-300>	120	Specifies the linger timer, which starts when a leave message is sent to the router. The multicast filter is deleted when the timer expires.
server-topology <centralized | distributed>	centralized	Specifies whether the server topology is centralized or distributed.
priority <NUMBER: 0-7>	7	Specifies the priority tag inserted into all IGMP messages.
query-delay <DECISECONDS: 1-100>	10	Specifies how long to wait after sending a query before sending the next. The delay must be greater than the query-response-interval.
query-engine <on | off>	off	Specifies whether the query engine is on or off.
query-interval <SECONDS: 10-31744>	125	Specifies how often to send the general query.
query-ip-source-addr <IP address>	0.0.0.0	Specifies the source address used for all general queries sent on an L2 forwarding domain. ***Note:*** This attribute is ignored when query-engine is disabled.
query-response-interval <DECISECONDS: 10-31744>	100	Specifies the amount of time to wait for a reply to the general query. This value goes directly into the query packet. This interval must be shorter than query-delay.
rapid-recovery-mode <on | off>	off	Specifies whether rapid recovery mode is on or off.
robustness <NUMBER: 1-2>	2	Specifies the number of query replies (IGMP Join) or Leaves to send for each multicast group.

Copyright© 2022 Ciena® Corporation

Table 9-19
Multicast-services attributes

Attribute	Default value	Description
router-address-range <IP multicast address range>	undefined, that is, 0.0.0.0	Specifies the range of IP addresses reserved for the router. This setting must remain undefined when the server-topology is set to distributed. Valid range is 224.0.0.3 - 239.255.255.255.
router-query-interval <SECONDS: 10-31744>	250	Used for selecting the IGMP querier when two or more routers exist on the same network.
Server Interface	empty	Is a list of interfaces.
UMF	flood	Used to specify an unknown multicast-group filter.

The procedures for configuring multicast services are:

- "Configuring a VLAN as a multicast L2 forwarding domain" on page 9-34
- "Configuring a multicast service on a virtual switch" on page 9-35
- "Configuring channel streams" on page 9-36
- "Configuring IGMP forking with VLAN translation" on page 9-38
- "Configuring a multicast router topology" on page 9-40
- "Configuring a multicast server topology" on page 9-44
- "Configuring multicast servers with redundant routers" on page 9-47
- "Configuring redundant query engines" on page 9-51
- "Clearing multicast service statistics" on page 9-55
- "Displaying multicast services information" on page 9-56
- "MPLS-based multicast configuration example" on page 9-58

Copyright© 2022 Ciena® Corporation

Procedure 9-1
Configuring a VLAN as a multicast L2 forwarding domain

Configure a VLAN to act as a multicast L2 forwarding domain for multicast services.

Step	Action
1	Enable multicast services: `multicast-services enable`
2	Enable multicast services on the switch: `multicast-services igmp-snooping enable`
3	Create a multi-service VLAN `multicast-service create vlan <vlan-id>`
4	Set attributes for the VLAN acting as a multicast L2 forwarding domain: `multicast-services igmp-snooping set vlan <vlan> {active-linger-timeout <SECONDS: 0-300>} {querier-compatibility-mode <v2|v3>} {default-router-port <Port>} {fork <on|off>} {last-member-query-interval <DECISECONDS: 10..100>} {leave-mode <fast|inquisitive>} {linger-timeout <SECONDS: 10-300>} {min-response-time <DECISECONDS: 50..600>} {server-topology <centralized|distributed>} {priority <NUMBER: 0-7>} {query-delay <DECISECONDS: 1..100>} {query-engine <on|off>} {query-interval <SECONDS: 10-31744>} {query-ip-source-addr <IP address>} {query-response-interval <DECISECONDS: 10..255>} {rapid-recovery-mode <on|off>} {robustness <NUMBER: 1-2>} {router-address-range <IP multicast address range>} {router-query-interval <SECONDS: 10-31744>} {querier-compatibility-mode <v2 |v3>} {server-topology <centralized | distributed>} {rapid-recovery-mode <on|off>}`

To add server ports to the multicast L2 forwarding domain

5	Add server ports to the multicast L2 forwarding domain: `multicast-services server-interface add vlan <vlan> [port <Port list>]`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 9-2
Configuring a multicast service on a virtual switch

Configure a multicast service on a virtual switch.

Step	Action
1	Enable multicast services: `multicast-services enable`
2	Create a multicast service on a virtual switch: `multicast-services create vs <vs>`
3	Enable IGMP snooping over a virtual switch: `multicast-services igmp-snooping enable vs <vs-object>`
4	Set attributes for the virtual acting as a multicast L2 forwarding domain: `multicast-services igmp-snooping set vs <vs> {active-linger-timeout <SECONDS: 0-300>} {default-router-port <Port>} {default-router-mpls-vc <mpls-vc-object-attached>} {leave-mode <fast|inquisitive>} {last-member-query-interval <DECISECONDS: 10-100>} {linger-timeout <SECONDS: 10-300>} {min-response-time <50-600 deciseconds>}{priority <NUMBER: 0-7>} {query-delay <DECISECONDS: 1-100>} {query-engine <on|off>} {query-interval <SECONDS: 10-31744>} {query-ip-source-addr <IP address>} {query-response-interval <DECISECONDS: 10-255>} {robustness <NUMBER: 1-2>}{router-address-range <IP multicast address range>} {router-query-interval <SECONDS: 10-31744>}{querier-compatibility-mode <v2|v3>} {server-topology <centralized | distributed>} {rapid-recovery-mode <on|off>}`

To add server ports to the multicast L2 forwarding domain

5	Add server ports to the multicast L2 forwarding domain: `multicast-services server-interface add vs <vs> [port <Port list>]`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 9-3
Configuring channel streams

Configure a channel stream to define a set of multicast groups that are flooded to all interfaces in the forwarding domain.

When multicast services are enabled for a particular forwarding domain, a new and empty channel stream is created for that forwarding domain. Multicast groups can be added or removed from this channel stream at any time. As multicast groups are added to the channel stream, they immediately begin to flow to all interfaces in the forwarding domain. Configured channel streams are saved in the system configuration.

You can:

- add a multicast group to a channel stream
- remove a multicast group from a channel stream
- exclude a interface from a channel stream
- include a interface in a channel stream

Step	Action

To add a multicast group to a channel stream

1 Add a multicast group to a channel stream:

```
multicast-services channel-stream add [vlan <vlan>| vs
<vs>] {group-address <IP multicast address range>}
```

where

vlan <vlan> | vs <vs>	is the VLAN or VS to add the multicast group for.
group-address <IP multicast address range>	is the IP multicast address or addresses for the multicast group.

Copyright© 2022 Ciena® Corporation

To remove a multicast group from a channel stream

2 Remove a multicast group from a channel stream:

```
multicast-services channel-stream remove [vlan <vlan>|vs
<vs>]{group-address <IP multicast address range>}
```

where

vlan <vlan> | vs <vs>	is the VLAN or VS to remove the channel stream group from.
group-address <IP multicast address range>	is the IP multicast address or addresses for the channel stream group.

To exclude a port from a channel stream

3 Exclude a port from a channel stream:

```
multicast-services channel-stream exclude [vlan <vlan> |
vs <vs>] {port <Port list>} mpls-vc <VcMplsAttached>
```

where

vlan <vlan> | vs <vs>	is the VLAN or VS.
port <Port list>	is the port or ports to exclude.
VcMplsAttached	is the MPLS virtual circuit to exclude.

To include a port in a channel stream

4 Include a port in a channel stream:

```
multicast-services channel-stream include [vlan <vlan>|vs
<vs>] {port <Port list>}mpls-vc <VcMplsAttached>
```

where

vlan <vlan> | vs <vs>	is the VLAN or VS.
port <Port list>	is the port or ports to include.
VcMplsAttached	is the MPLS virtual circuit to include.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 9-4
Configuring IGMP forking with VLAN translation

To configure IGMP forking with VLAN translation, IGMP snooping must be enabled.

Step	Action
1	Enable IGMP snooping: `multicast-services igmp-snooping enable`
2	Create a multicast VLAN: `vlan create vlan <VLAN>`
3	Add a host port: `vlan add vlan <VLAN> {port <Port>}`
4	Add a router port: `vlan add vlan <VLAN> {port <Port>}`
5	Create a multicast service: `multicast-services create vlan <VLAN>`
6	Turn IGMP forking on: `multicast-services igmp-snooping set vlan <VLAN> fork on`
7	Create the upstream data vlan or virtual switch: `vlan create vlan <VLAN>`
8	Add the upstream data VLAN or to the host and router port: `vlan add {vlan <VLAN>} {port <Port>}` `vlan add {vlan <VLAN>} {port <Port>}`
9	Configure VLAN translation to the host port: `vlan translate add {port <Port>} vid <NUMBER: 0-4095> vlan {vlan <VLAN>}`

—end—

Copyright© 2022 Ciena® Corporation

Example

This example shows IGMP forking with VLAN translation. This example assumes that your multicast service is on vlan 100 and the upstream data vlan is 101.

```
vlan create vlan 100
vlan add vlan 100 port 1
vlan add vlan 100 port 3
multicast-services create vlan 100
multicast-services igmp-snooping set vlan 100 fork on
vlan create vlan 101
vlan add vlan 101 port 1
vlan add vlan 101 port 3
vlan translate add port 1 vid 100 vlan 101
```


Copyright© 2022 Ciena® Corporation

Procedure 9-5
Configuring a multicast router topology

Configure devices in a network where a multicast server is delivering services over the network by means of a multicast router, where the router is also the querier and controls the delivery of multicast streams.

In this example, multicast services are delivered on VLAN 300.

Figure 9-10
Multicast services with a multicast router

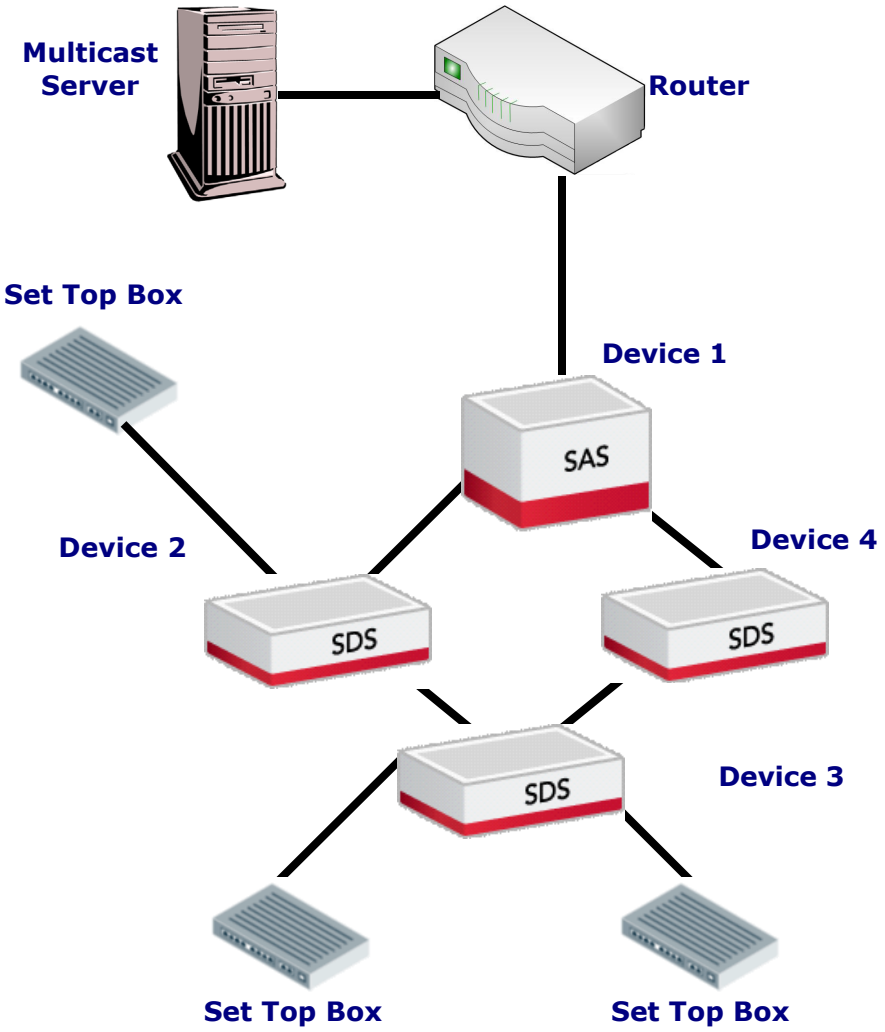

The default value for the query-ip-source-address is 0.0.0.0. This value is used as the source address for all query messages as well as for proxy-reply messages sent to the router. The default value works in most cases. If this

Copyright© 2022 Ciena® Corporation

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.

For rapid recovery queries IP SA 0.0.0.0 is only used. No other queries with IP SA=0.0.0.0 are sent out. If query-engine is enabled, IP SA must be set to a value other than 0.0.0.0. Otherwise, query-engine cannot be enabled. Membership reports with IP SA 0.0.0.0 are sent.

For proxy query, the IP SA of the router that sent the query is used.

Step	Action
1	Create a multicast entry for the VLAN on all devices, which enables multicast services for the specified VLAN: `multicast-services create vlan <VlanList>`
2	Set attributes for igmp-snooping. Note that query timers must be set to the same values as on the multicast router: `multicast-services igmp-snooping set vlan <vlan> {active-linger-timeout <SECONDS: 0-300>} {querier-compatibility-mode <v2|v3>} {default-router-port <Port>} {fork <on|off>} {last-member-query-interval <DECISECONDS: 10..100>} {leave-mode <fast|inquisitive>} {linger-timeout <SECONDS: 10-300>} {min-response-time <DECISECONDS: 50..600>} {server-topology <centralized|distributed>} {priority <NUMBER: 0-7>} {query-delay <DECISECONDS: 1..100>} {query-engine <on|off>} {query-interval <SECONDS: 10-31744>} {query-ip-source-addr <IP address>} {query-response-interval <DECISECONDS: 10..255>} {rapid-recovery-mode <on|off>} {robustness <NUMBER: 1-2>} {router-address-range <IP multicast address range>} {router-query-interval <SECONDS: 10-31744>} {querier-compatibility-mode <v2 |v3>} {server-topology <centralized | distributed>} {rapid-recovery-mode <on|off>}`
3	Enable IGMP snooping: `multicast-services igmp-snooping enable`
4	To enable UMF on the specified VLAN for the devices that have subscriber equipment connected to them, enter this command: `multicast-services umf drop vlan <vlan> | vs <vs>` 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.

Copyright© 2022 Ciena® Corporation

5 To verify the IGMP settings, enter this command:

```
multicast-services show configuration
```

—end—

Example

This example configures device 2 as seen in Figure 9-10.

```
multicast-services create vlan 986

multicast-services igmp-snooping set vlan 986 active-linger-
timeout 30 querier-compatibility-mode v2 default-router-port 0
fork off last-member-query-interval 10 leave-mode fast linger-
timeout 125 min-response-time 50 server-topology centralized
priority 7 query-delay 10 query-engine off query-interval 125
query-response-interval 100 rapid-recovery-mode off robustness
1 router-address-range 0.0.0.0 router-query-interval 125

multicast-services igmp-snooping enable
multicast-services umf drop vlan 300
multicast-services show configuration vlan 300
+-------------------------------------------------------------------+
| Multicast-Services: Global Enable                                 |
| IGMP-Snooping: Global Enable                                      |
+-----------------------------------------------+-------------------+

+-----------------------------------------------+-------------------+
|    MULTICAST CONFIGURATION VLAN 300 VLAN#300                      |
+-----------------------------------------------+-------------------+
| Multicast Services Admin State                | enable            |
| UMF                                           | drop              |
| WKM Forwarding                                | Disabled          |
+-----------------------------------------------+-------------------+
|    IGMP CONFIGURATION (general)                                   |
+-----------------------------------------------+-------------------+
| IGMP Snooping Admin State                     | enable            |
| IGMP Forking                                  | off               |
| IGMP Server Topology                          | centralized       |
| IGMP Query Engine                             | off               |
| IGMP Leave Mode                               | fast              |
| IGMP Querier Compatibility Mode               | v2                |
| L2 Packet Priority                            | 7                 |
| Robustness                                    | 2                 |
+-----------------------------------------------+-------------------+
|    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                |
```


Copyright© 2022 Ciena® Corporation

```
| 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                   |
| Rapid Query Mode                              | Off                 |
+-----------------------------------------------+---------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 9-6
Configuring a multicast server topology

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.

Step	Action
1	Create a multicast entry for the VLAN or virtual switch on all devices: `multicast-services create [vlan <VlanList>|vs <vs>]` This enables multicast services for the specified VLAN.
2	Configure the query engine on the same device: `multicast-services igmp-snooping set [vlan <vlan>|vs <vs>] {active-linger-timeout <SECONDS: 0-300>} {querier-compatibility-mode <v2|v3>} {default-router-port <Port>} {fork <on|off>} {last-member-query-interval <DECISECONDS: 10..100>} {leave-mode <fast|inquisitive>} {linger-timeout <SECONDS: 10-300>} {min-response-time <DECISECONDS: 50..600>} {server-topology <centralized|distributed>} {priority <NUMBER: 0-7>} {query-delay <DECISECONDS: 1..100>} query-engine on {query-interval <SECONDS: 10-31744>} query-ip-source-addr 10.10.10.10 {query-response-interval <DECISECONDS: 10..255>} {rapid-recovery-mode <on|off>} {robustness <NUMBER: 1-2>} {router-address-range <IP multicast address range>} {router-query-interval <SECONDS: 10-31744>}`
3	Enable IGMP snooping: `multicast-services igmp-snooping enable`
4	To enable UMF on the specified VLAN or virtual switch for the devices that have subscriber equipment connected to them, enter this command: `multicast-services umf drop vlan <VlanList>|vs <vs>`
5	To verify the IGMP settings, enter this command: `multicast-services show configuration`

—end—

Copyright© 2022 Ciena® Corporation

Example

For this example, multicast services are delivered on VLAN 300. In the topology represented in the Delivering multicast services using the query engine figure, UMF is required on device 1 to contain the streams being flooded by the multicast server. On all other devices, UMF is optional.

Figure 9-11
Delivering multicast services using the query engine

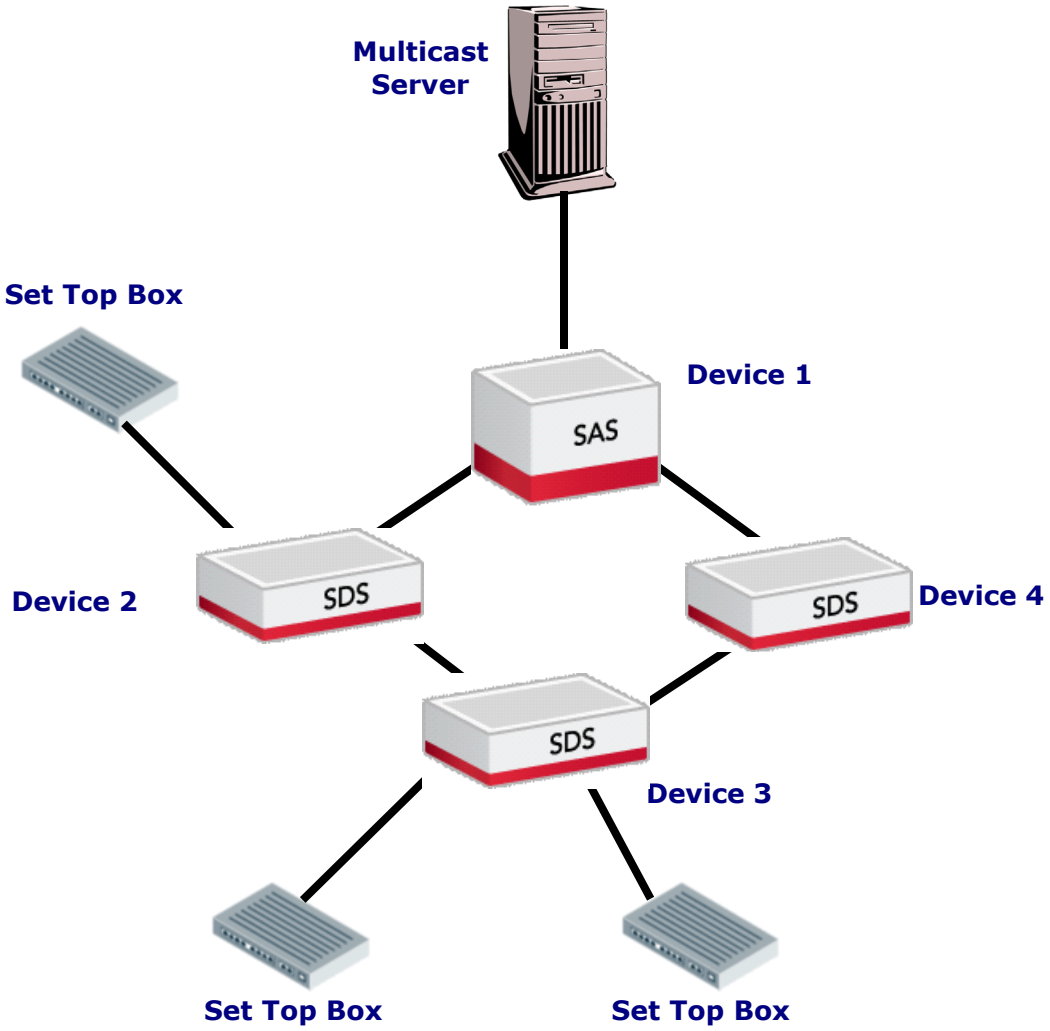

This example configures device 1 as shown in Figure 9-11.

```
multicast-services create vlan 300
```


Copyright© 2022 Ciena® Corporation

```
              multicast-services igmp-snooping set vlan 300 active-linger-
              timeout 30 querier-compatibility-mode v2 default-router-port 0
              fork off last-member-query-interval 10 leave-mode fast linger-
              timeout 125 min-response-time 50 server-topology centralized
              priority 7 query-delay 10 query-interval 125 query-ip-source-
              addr 10.10.10.10 query-engine on query-response-interval 100
              rapid-recovery-mode off robustness 1 router-address-range
              0.0.0.0 router-query-interval 125

              multicast-services igmp-snooping enable
              multicast-services umf drop vlan 300
              multicast-services show configuration vlan 300
+-------------------------------------------------------------------------+
| Multicast-Services: Global Enable                                       |
| IGMP-Snooping: Global Enable                                            |
+-----------------------------------------------+---------------------+

+-----------------------------------------------+---------------------+
|    MULTICAST CONFIGURATION VLAN 300 VLAN#300                            |
+-----------------------------------------------+---------------------+
| Multicast Services Admin State                | enable              |
| UMF                                           | drop                |
| WKM Forwarding                                | Disabled            |
+-----------------------------------------------+---------------------+
|    IGMP CONFIGURATION (general)                                         |
+-----------------------------------------------+---------------------+
| IGMP Snooping Admin State                     | enable              |
| IGMP Forking                                  | off                 |
| IGMP Server Topology                          | centralized         |
| IGMP Query Engine                             | on                  |
| IGMP Leave Mode                               | fast                |
| IGMP Querier Compatibility Mode               | v2                  |
| L2 Packet Priority                            | 7                   |
| Robustness                                    | 2                   |
+-----------------------------------------------+---------------------+
|    IGMP CONFIGURATION (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 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                   |
| Rapid Query Mode                              | Off                 |
+-----------------------------------------------+---------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 9-7
Configuring multicast servers with redundant routers

Configure devices in a network that has two or more redundant multicast servers. Both multicast servers are connected to the network by means of a multicast router. In this scenario, only one router, with its associated multicast server, is active at a time. The second router and server are in a standby mode. The active server is determined by identifying the router that has the lowest IP address. This router 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 the device goes down, the device advertises 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 send out general queries advertising a router address of 0.0.0.0. Device 4 becomes the interface for multicast services and all joins/leaves are then sent to it. When Router 1 is restored, it again becomes the primary router and Router 2 again goes 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.

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 (when Query-Engine is on and active), as well as for proxy-reply messages sent to the router. Without query-engine on, the default value works in most cases. With query-engine on, the query-ip source-addr must be set to a valid IP other than 0.0.0.0. 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 or virtual switch on all devices: `multicast-services create [vlan <VlanList>|vs <vs>]` This enables multicast services for the specified VLAN.

Copyright© 2022 Ciena® Corporation

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 vlan <VlanList> | vs <vs>
```

3 Set attributes for igmp-snooping. Note that query timers must be set to the same values as on the multicast router:

```
multicast-services igmp-snooping set [vlan <vlan>|vs
<vs>] {active-linger-timeout <SECONDS: 0-300>} {querier-
compatibility-mode <v2|v3>} {default-router-port <Port>}
{fork <on|off>} {last-member-query-interval
<DECISECONDS: 10..100>} {leave-mode <fast|inquisitive>}
{linger-timeout <SECONDS: 10-300>} {min-response-time
<DECISECONDS: 50..600>} {server-topology
<centralized|distributed>} {priority <NUMBER: 0-7>}
{query-delay <DECISECONDS: 1..100>} {query-engine
<on|off>} {query-interval <SECONDS: 10-31744>} {query-ip-
source-addr <IP address>} {query-response-interval
<DECISECONDS: 10..255>} {rapid-recovery-mode <on|off>}
{robustness <NUMBER: 1-2>} {router-address-range <IP
multicast address range>} {router-query-interval
<SECONDS: 10-31744>}
```

4 Enable IGMP snooping on all four devices.

```
multicast-services igmp-snooping enable
```

5 (Optional) Verify the IGMP settings.

```
multicast-services show configuration
```

—end—

Example

In this example, multicast services are assumed to be delivered on VLAN 300. Figure 9-12 shows a sample network topology for delivering multicast services using redundant multicast routers. For the network represented in Figure 9-12, UMF is required on device 1, device 2 and device 4. On all other devices, it is optional.

Copyright© 2022 Ciena® Corporation

Figure 9-12
Delivering multicast services using redundant multicast routers

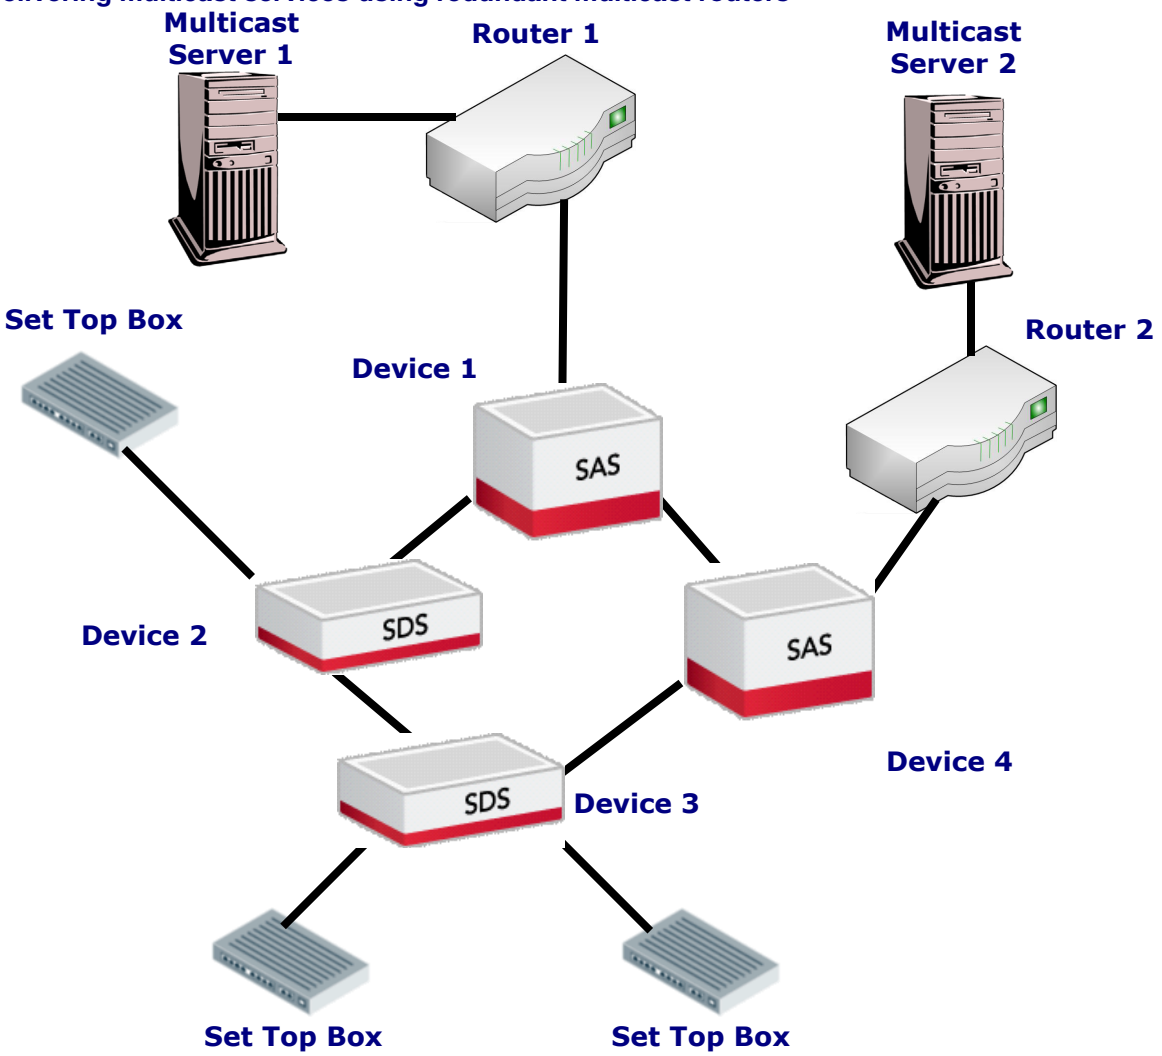

This example configures device 1 as shown in the Delivering multicast services using redundant multicast routers figure. Note that device 1 and device 4 have the same configuration.

```
multicast-services create vlan 300
multicast-services umf drop vlan 300
multicast-services igmp-snooping set vlan 300 active-linger-
timeout 30 querier-compatibility-mode v2 default-router-port 0
fork off last-member-query-interval 10 leave-mode fast linger-
timeout 125 min-response-time 50 server-topology centralized
priority 7 query-delay 10 query-engine off query-interval 125
```


Copyright© 2022 Ciena® Corporation

```
query-response-interval 100 rapid-recovery-mode off robustness
1 router-address-range 0.0.0.0 router-query-interval 125
multicast-services igmp-snooping enable
multicast-services show configuration vlan 300
```

```
+-------------------------------------------------------------------------+
| Multicast-Services: Global Enable                                       |
| IGMP-Snooping: Global Enable                                            |
+-----------------------------------------------------+---------------------+

+-----------------------------------------------------+---------------------+
|    MULTICAST CONFIGURATION VLAN 300 VLAN#300                              |
+-----------------------------------------------------+---------------------+
| Multicast Services Admin State                      | enable              |
| UMF                                                 | drop                |
| WKM Forwarding                                      | Disabled            |
+-----------------------------------------------------+---------------------+
|    IGMP CONFIGURATION (general)                                           |
+-----------------------------------------------------+---------------------+
| IGMP Snooping Admin State                           | enable              |
| IGMP Forking                                        | off                 |
| IGMP Server Topology                                | centralized         |
| IGMP Query Engine                                   | off                 |
| IGMP Leave Mode                                     | fast                |
| IGMP Querier Compatibility Mode                     | v2                  |
| L2 Packet Priority                                  | 7                   |
| Robustness                                          | 2                   |
+-----------------------------------------------------+---------------------+
|    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                   |
| Rapid Query Mode                                    | Off                 |
+-----------------------------------------------------+---------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 9-8
Configuring redundant query engines

Configure redundant query engines.

Step	Action
1	Create a multicast entry for the VLANs or virtual switches: `multicast-services create [vlan <VlanList>|vs <vs>}` This enables multicast services for the specified VLAN.
2	Configure the query engine on devices 1 and 3: `multicast-services igmp-snooping set [vlan <vlan>|vs <vs>]{active-linger-timeout <SECONDS: 0-300>} {querier-compatibility-mode <v2|v3>} {default-router-port <Port>} {fork <on|off>} {last-member-query-interval <DECISECONDS: 10..100>} {leave-mode <fast|inquisitive>} {linger-timeout <SECONDS: 10-300>} {min-response-time <DECISECONDS: 50..600>} {server-topology <centralized|distributed>} {priority <NUMBER: 0-7>} {query-delay <DECISECONDS: 1..100>} {query-engine <on|off>} {query-interval <SECONDS: 10-31744>} {query-ip-source-addr <IP address>} {query-response-interval <DECISECONDS: 10..255>} {rapid-recovery-mode <on|off>} {robustness <NUMBER: 1-2>} {router-address-range <IP multicast address range>} {router-query-interval <SECONDS: 10-31744>}` When using this command, the source IP address must come before the Query Engine is enabled, to function properly. Make sure that the source IP address for the two devices is different and that the device that becomes the primary query engine has a lower IP address.
3	Enable IGMP snooping on devices 1, 2, and 3: `multicast-services igmp-snooping enable`
4	Enable UMF on devices that have subscriber equipment connected to them: `multicast-services umf drop vlan <VlanList>|vs <vs>`
5	(Optional) Verify the IGMP settings. `multicast-services show configuration`

—end—

Copyright© 2022 Ciena® Corporation

Example

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 9-13
IGMP Snooping

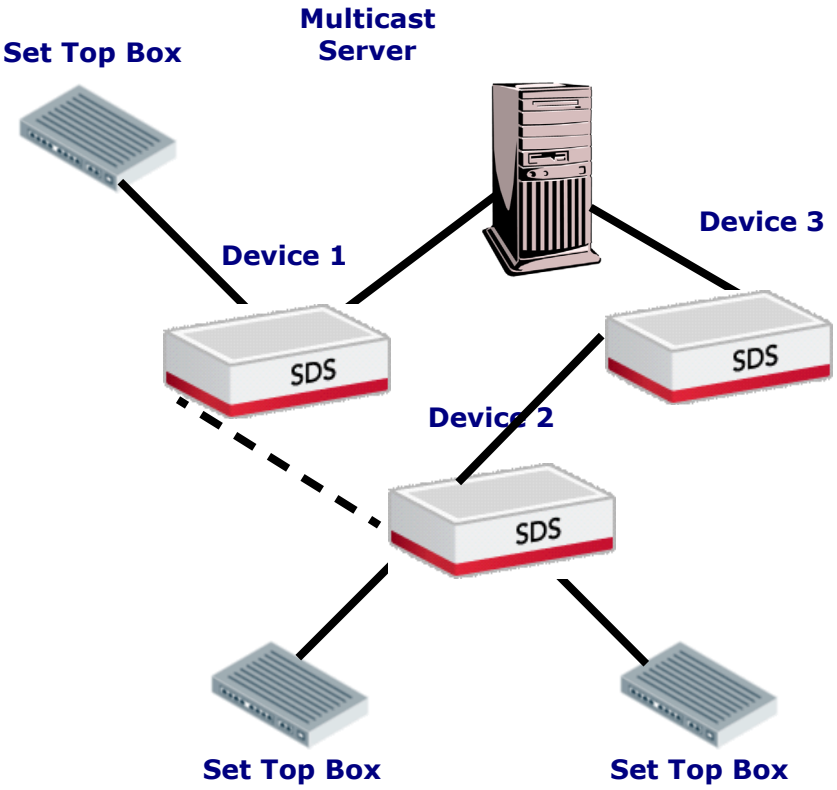

In this scenario, the query engine must be enabled on device 1 and device 3, but only one query engine is active at a time. The second query engine is in standby mode. The active query engine is determined by identifying the engine with the lowest source IP address. This query engine then handles general 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 advertises general query with a router source address of 0.0.0.0 to indicate that it has lost the router.

For example, assuming that device 1 has a lower source IP address than device 3, if the link between device 1 and device 2 goes down, device 2 sends out general queries with a router source address of 0.0.0.0 (rapid-recovery).

Copyright© 2022 Ciena® Corporation

The query-engine on device 3 then goes active and sends queries to device 1 to bring up its router interface. All membership reports from device 2 are then sent to device 3. When the link between device 1 and device 2 is restored, device 1 becomes the primary query-engine again, and device 3 returns to the standby state.

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.

For the network represented in Figure 9-13, multicast services are assumed to be delivered on VLAN 300. UMF is required on device 1 and device 3. On all other devices, it is optional. The network is using RSTP.

This example configures device 1 as shown in Figure 9-11. 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 vlan 300 active-linger-
timeout 30 querier-compatibility-mode v2 default-router-port 0
fork off last-member-query-interval 10 leave-mode fast linger-
timeout 125 min-response-time 50 server-topology centralized
priority 7 query-delay 10 query-interval 125 query-ip-source-
addr 10.10.10.10 query-engine on query-response-interval 100
rapid-recovery-mode off robustness 1 router-address-range
0.0.0.0 router-query-interval 125
multicast-services igmp-snooping enable
multicast-services umf drop vlan 300
multicast-services show configuration vlan 300
```


Copyright© 2022 Ciena® Corporation

```
+-------------------------------------------------------------------------+
| Multicast-Services: Global Enable                                       |
| IGMP-Snooping: Global Enable                                            |
+-----------------------------------------------------+---------------------+

+-----------------------------------------------------+---------------------+
|    MULTICAST CONFIGURATION VLAN 300 VLAN#300                              |
+-----------------------------------------------------+---------------------+
| Multicast Services Admin State                      | enable              |
| UMF                                                 | drop                |
| WKM Forwarding                                      | Disabled            |
+-----------------------------------------------------+---------------------+
|    IGMP CONFIGURATION (general)                                           |
+-----------------------------------------------------+---------------------+
| IGMP Snooping Admin State                           | enable              |
| IGMP Forking                                        | off                 |
| IGMP Server Topology                                | centralized         |
| IGMP Query Engine                                   | on                  |
| IGMP Leave Mode                                     | fast                |
| IGMP Querier Compatibility Mode                     | v2                  |
| L2 Packet Priority                                  | 7                   |
| Robustness                                          | 2                   |
+-----------------------------------------------------+---------------------+
|    IGMP CONFIGURATION (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 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                   |
| Rapid Query Mode                                    | Off                 |
+-----------------------------------------------------+---------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 9-9
Clearing multicast service statistics

You can clear:

- multicast services statistics
- multicast services statistics by VLAN
- multicast service statistics by virtual switch

Step	Action

To clear multicast services statistics

1 Clear multicast services statistics

```
multicast-services clear {statistics}
```

To clear multicast services statistics by VLAN

2 Clear multicast services statistics by VLAN

```
multicast-services clear vlan <vlan> {statistics}
```

To clear multicast services statistics by virtual switch

3 Clear multicast services statistics by virtual switch

```
multicast-services clear vs <vs> {statistics}
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 9-10
Displaying multicast services information

Display multicast services information to verify the configuration.

You can display:

- a brief summary of multicast configuration and status
- all configuration settings
- all configuration and status
- group status
- active multicast groups
- all configuration settings

Step	Action

To display a brief summary of multicast configuration and status

1 Display a brief summary of multicast configuration and status:

```
multicast-services show brief
```

To display all configuration settings

2 Display all configuration settings:

```
multicast-services show configuration {vlan <VLAN
list>}{vs <vs>}
```

To display all configuration and status

3 Display all configuration and status:

```
multicast-services show everything {vlan <VLAN list>}{vs
<vs>}
```

To display group status

4 Display group status:

```
multicast-services show group-status {vlan <VLAN
list>}{vs <vs>}{group-address} <IP multicast address
range>}
```

To display all active multicast groups

5 Display all active multicast groups:

```
multicast-services show groups {vlan <VLAN list>} {vs
<vs>} {port <Port list>}{mpls-vc}
```


Copyright© 2022 Ciena® Corporation

To display all configuration settings

6 Display all configuration settings:

```
multicast-services show status {vlan <VLAN list>}{vs
<vs>}
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 9-11
MPLS-based multicast configuration example

This figure shows a configuration example for MPLS-based multicast. The configuration steps follow.

Figure 9-14
MPLS-based multicast configuration example

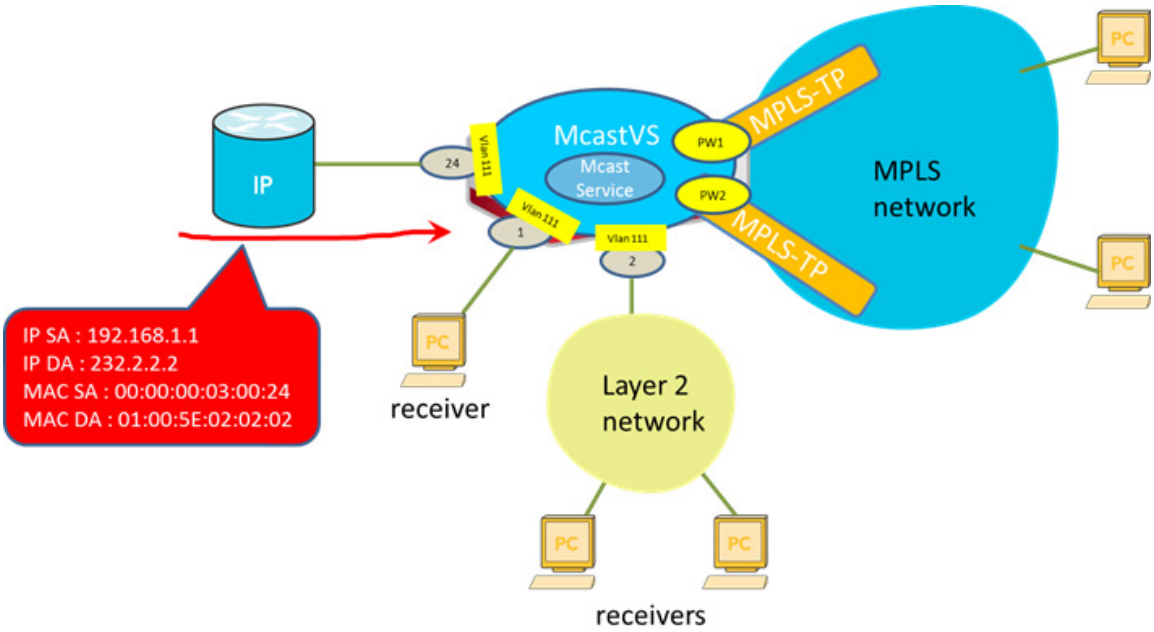

Step	Action
1	Create the virtual switch MCastVS: `virtual-switch ethernet add vs MCastVS mode vpls`
2	Add the virtual switch MCastVS to port 1 and vlan 111: `virtual-switch ethernet add vs MCastVS port 1 vlan 111 ingress-l2-transform pop egress-l2-transform push-8100.111`
3	Add the virtual switch MCastVS to port 2 and vlan 111: `virtual-switch ethernet add vs MCastVS port 2 vlan 111 ingress-l2-transform pop egress-l2-transform push-8100.111`
4	Add the virtual switch MCastVS to port 24 and vlan 111: `virtual-switch ethernet add vs MCastVS port 24 vlan 111 ingress-l2-transform pop egress-l2-transform push-8100.111`

Copyright© 2022 Ciena® Corporation

5 Attach the virtual circuit PW1 to the MPLS virtual switch MCastVS:

```
virtual-switch ethernet attach mpls-vc PW1 vs MCast VS
```

6 Attach the virtual circuit PW2 to the MPLS virtual switch MCastVS:

```
virtual-switch ethernet attach mpls-vc PW2 vs MCast VS
```

7 Create a multicast service on the virtual switch MCastVS:

```
multicast-services create vs MCastVS
```

8 Enable UMF on the virtual switch MCastVS:

```
multicast-services umf drop vs MCastVS
```

9 Specify the source address of 10.10.10.33 to use for all general queries:

```
multicast-services igmp-snooping set vs MCastVS query-ip-
source-addr 10.10.10.33
```

10 Enable IGMP snooping:

```
multicast-services igmp-snooping enable
```

—end—

Example

This example shows the output from the multicast-services show everything command for vsMcastVS:

```
> multicast-services show everything vs MCastVS
+-------------------------------------------------------------------------------+
| Multicast-Services Global State : Enable                                      |
| IGMP-Snooping Global Admin State: Enable                                      |
| IGMP-Snooping Global Oper State : Enable                                      |
+-------------------------------------------------------------------------------+

+-----------------------------------------------------+------------------------+
|    MULTICAST CONFIGURATION   vs MCastVS                                      |
+-----------------------------------------------------+------------------------+
| Multicast Services Admin State                      | enable                 |
| UMF                                                 | drop                   |
| WKM Forwarding                                      | Disabled               |
+-----------------------------------------------------+------------------------+
|    IGMP CONFIGURATION (general)                                              |
+-----------------------------------------------------+------------------------+
| IGMP Snooping Admin State                           | enable                 |
| IGMP Forking                                        | off                    |
| IGMP Server Topology                                | centralized            |
| IGMP Query Engine                                   | off                    |
| IGMP Leave Mode                                     | fast                   |
| IGMP Querier Compatibility Mode                     | v3                     |
| L2 Packet Priority                                  | 7                      |
| Robustness                                          | 2                      |
+-----------------------------------------------------+------------------------+
|    IGMP CONFIGURATION (proxy query)                                          |
+-----------------------------------------------------+------------------------+
| Proxy Query Interval (s)                            | 125                    |
| Proxy Query Response Interval (ds)                  | 100                    |
| Proxy Query Delay (ds)                              | 10                     |
| Proxy Query Source Address                          | 10.10.10.33            |
| Last Member Query Interval (ds)                     | 10                     |
+-----------------------------------------------------+------------------------+
|    IGMP CONFIGURATION (router)                                               |
+-----------------------------------------------------+------------------------+
| Router Query Interval (s)                           | 250                    |
```


Copyright© 2022 Ciena® Corporation

```
| 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 Interface                              | 0                       |
| Rapid Query Mode                                      | Off                     |
+-------------------------------------------------------+-------------------------+
+-------------------------------------------------------+-------------------------+
|    IGMP Operational Status    vs MCastVS                                          |
+-------------------------------------------------------+-------------------------+
| Multicast services Oper State                         | Enable                  |
| IGMP Snooping Oper State                              | Enable                  |
| Query Engine                                          | Disabled                |
| Next Proxy Query (s)                                  | 30                      |
| Sending Reports to                                    | router interface        |
+-------------------------------------------------------+-------------------------+
|    EXTERNAL IGMP QUERIER (router)                     |                         |
+-------------------------------------------------------+-------------------------+
| MAC address                                           | 00:00:00:03:00:24       |
| IP address                                            | 10.10.10.24             |
| IGMP Version                                          | v3                      |
| Interface                                             | port 24                 |
| Mesh VC                                               | No                      |
| Older Version Querier Time remaining                  | n/a                     |
| Query interval time remaining                         | 239 seconds             |
+-------------------------------------------------------+-------------------------+
|    MULTICAST GROUP STATUS                                                         |
+-------------------------------------------------------+-------------------------+
| Channel Stream Groups                                 | 0                       |
| Dynamic Groups                                        | 1                       |
| Active groups from the router                         | 1                       |
| Lingering groups                                      | 0                       |
| Total dynamic group members                           | 4                       |
+-------------------------------------------------------+-------------------------+
|   IGMP STATISTICS                                                                 |
+-------------------------------------------------------+-------------------------+
| Membership Reports Rx                                 | 0                       |
| Leave Reports Rx                                      | 0                       |
| V3 membership Reports Rx                              | 24086                   |
|   IS_IN Reports Rx                                    | 23369                   |
|   IS_EX Reports Rx                                    | 692                     |
|   TO_IN Reports Rx                                    | 2                       |
|   TO_EX Reports Rx                                    | 2                       |
|   Allow Reports Rx                                    | 8                       |
|   Block Reports Rx                                    | 4                       |
| Queries Received                                      | 1504                    |
| Query Rx Discards                                     | 4                       |
| Router Timeouts                                       | 3                       |
| Query Reply Timeouts                                  | 0                       |
| Router Rx Discards                                    | 0                       |
| Host Rx Discards                                      | 0                       |
| Checksum Errors Rx                                    | 0                       |
| L2/L3 Mismatch Rx                                     | 0                       |
| Unknown Packet Type Rx                                | 0                       |
| Exceeded Resources                                    | 0                       |
| Query Discarded because of rapid recovery mode off    | 0                       |
| Compatibility Mode Discards                           | 0                       |
+-------------------------------------------------------+-------------------------+
+----------+----------+-----+----+------+------+--------------+-------+--------------+-------------+
| Service  |IP address|state|type|source|member| uptime       |timeout| compatibility| group-timer |
|          |          |     |    |      |count |              |       | mode         |             |
+----------+----------+-----+----+------+------+--------------+-------+--------------+-------------+
|vs MCastVS| 232.2.2.2| act |dyna|router|   4  | 0d 5h 42m 12s|       |           v3 |  177 seconds|
+----------+----------+-----+----+------+------+--------------+-------+--------------+-------------+
```


Copyright© 2022 Ciena® Corporation

```
S3*> multicast-services show group-status group-address 232.2.2.2

+-------------+----------+-----+----+------+------+--------------+-------+---------------+------------+
| Service     |IP address|state|type|source|member| uptime       |timeout| compatibility-| group-timer|
|             |          |     |    |      |count |              |       | mode          |            |
+-------------+----------+-----+----+------+------+--------------+-------+---------------+------------+
| vs   MCastVS| 232.2.2.2| act |dyna|router|    4 | 0d 5h 46m 34s|       |         v3    | 166 seconds|
+-----------------------------------------------------------------------------------------------------+
|     IGMP V2 Host Present Time remaining|                                                        N/A  |
+-----------------------------------------------------------------------------------------------------+
| GROUP MEMBER Ports: 1 2                                                                             |
| GROUP MEMBER mpls-vcs: PW1  PW2                                                                     |
+-----------------------------------------------------------------------------------------------------+
| Group Source Records:                                                                               |
| Group Filter mode: EXCLUDE                                                                          |
| Number of source addresses:  2                                                                      |
|                                     Source Address Records                                          |
+-----------------------------------------------------------------------------------------------------+
|     Source Address                          |        Source Timeout expected                        |
+---------------------------------------------+-------------------------------------------------------+
|                            192.168.1.1      |                                          259 seconds  |
|                            192.168.1.2      |                                          259 seconds  |
+-----------------------------------------------------------------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Copyright© 2022 Ciena® Corporation

TDM services over PSN configuration

Time Division Multiplexing (TDM) services allow the network provider to create TDM connections through a Packet Switched Network (PSN). TDM services are available on these switches:

- 3903
- 3903x
- 3904
- 3905
- 3906
- 3924
- 3926
- 3928
- 5142
- 5160

TDM SFP ports on these switches provide TDM services. Additionally, TDM services are provided by

- ports on the Combo TDM OC48 FRU module in the 3926
- ports on the TDM FRU module in the 3926

TDM services are:

- Structure-Agnostic TDM over Packet (SAToP)
- Circuit Emulation Services over Packet Switched Networks (CESoP)

SAToP TDM pseudowire takes the raw TDM bit-stream from T1 or E1, disregards any TDM structure imposed on it, prepends the pseudowire header and dispatches the resulting packet stream over the Ethernet or MPLS PSN. At the remote end, the pseudowire header information is used to identify the egress attachment circuit. The TDM payload from the pseudowire is played out over the egress attachment circuit using the clock information as defined when the pseudowire is configured.

Copyright© 2022 Ciena® Corporation

With CESoP, structured TDM signals are encapsulated in a pseudowire and sent across the PSN. CESoP services are supported on TDM SFP ports. Advantages of CESoP are:

- PSN bandwidth savings
- DS0 level grooming
- better resiliency to packet loss in the PSN
- more flexible distribution

Bulk SFP upgrade command allows bulk upgrades of all SFPs on a device. A global default customer TPID (C-TPID), port default C-TPID and per pseudowire C-TPID is supported on the 3926 ports and virtual circuits attached to the 3926 ports. The global default applies to both module ports and SFP ports. Signal degrade and signal fail are also supported for TDM SFPs. TDM SFPs may override the global default C-TPID.

This section provides these topics:

- “Ports” on page 10-2
- “Signal degrade and signal fail” on page 10-7
- “Protection groups” on page 10-8
- “Encapsulation” on page 10-9
- “Transport” on page 10-10
- “E1 channel number mapping” on page 10-11
- “Linear APS 1+1” on page 10-17
- “Channel Associated Signaling (CAS) for nxDSO service” on page 10-31

Ports

TDM services are provided by:

- “TDM SFP ports” on page 10-3
- “TDM module ports” on page 10-6
- “Combo TDM OC48 FRU module ports” on page 10-6

Copyright© 2022 Ciena® Corporation

TDM SFP ports

This table summarizes TDM SFP port information.

Table 10-1
TDM SFP port information

TDM port type	Services supported	Loopback channelization	Supported channels	Attachment circuit channelization	TDM mode
OC3-STM1	CESoP	Port		STS3c	ANSI
OC3-STM1	CESoP	STS1	0, 1, 2	STS1	ANSI
OC3-STM1	CESoP	STS3c	0	STS3c	ANSI
OC12-STM4	CESoP	Port		STS12c	ANSI
OC12-STM4	CESoP	STS1	0 through 11	STS3c, Sts1	ANSI
OC12-STM4	CESoP	STS3c	0, 1, 2, 3 ***Note:*** • 0 overlaps with STS1 channels 0 to 2 • 1 overlaps with STS1 channels 3 to 5 • 2 overlaps with STS1 channels 6 to 8 • 3 overlaps with STS1 channels 9 to 11	STS3c, Sts1	ANSI
OC12-STM4	CESoP	STS12c	0	STS12c	ANSI
DS3-E3	CESoP	Port		STS1	ANSI
DS3-E3	CESoP	STS1	0	STS1	ANSI
OC3-STM1	SAToP	E1	0-62	E1	ETSI
OC3-STM1	SAToP	T1	0-83	T1	ANSI
OC3-STM1	CESoP	VC4c	0	VC4c	ETSI
OC12-STM4	CESoP	VC4c VC4-4c	0, 1, 2, 3 0	VC4c VC4-4c	ETSI ETSI
OC48-STM16	SAToP	Port	0	None	ANSI / ETSI
E1	SAToP	E1	0	E1	ETSI
T1	SAToP	T1	0	T1	ANSI

Copyright© 2022 Ciena® Corporation

Table 10-1
TDM SFP port information

TDM port type	Services supported	Loopback channelization	Supported channels	Attachment circuit channelization	TDM mode
OC48-STM16-CEM (Combo TDM OC48 FRU)	CESoP	STS48	0	STS48	ANSI
	CESoP	STS12c, STS3c, STS1	STS12c: 0-3 STS3c: 0-15 STS1: 0-47	STS12c, STS3c, STS1	ANSI
	SAToP	DS1	DS1: 0-1343	DS1 (See Note 2.)	ANSI
	SAToP	E1	E1: 0-1007	E1	ETSI
	CESoP	VC4-16c, VC4	VC4-16c: 0 VC4: 0-15	VC4-16c, VC4	ETSI
OC12-STM4-CEM (Combo TDM OC48 FRU)	CESoP	STS12c	0	STS12	ANSI
	CESoP	STS3c, STS1	STS3c: 0-3 STS1: 0-11	STS3c, STS1	ANSI
	SAToP	DS1	0-335	DS1 (See Note 2.)	ANSI
	SAToP	E1	E1: 0-251	E1	ETSI
	CESoP	VC4-4c, VC4	VC4-4c: 0 VC4: 0-3	VC4-4c, VC4	ETSI
OC3-STM1-CEM (Combo TDM OC48 FRU)	CESoP	STS3c	0	STS3c	ANSI
	CESoP	STS1	0-2	STS1	ANSI
	SAToP	DS1	0-83	DS1 (See Note 2.)	ANSI
	SAToP	E1	E1: 0-62	E1	ETSI
	CESoP	VC4	VC4: 0	VC4	ETSI
DS3-E3 (Combo TDM OC48 FRU)	SAToP	Port, DS1 (See Note 4.)	Port: 0 DS1: 0-27	DS3, DS1 (See Note 1.) (See Note 2.)	ANSI
	CESoP	STS1	0	STS1	ANSI
DS3-E3 (Combo TDM OC48 FRU)	SAToP	Port (See Note 5.)	0	E3	ETSI

Copyright© 2022 Ciena® Corporation

Table 10-1
TDM SFP port information

TDM port type	Services supported	Loopback channelization	Supported channels	Attachment circuit channelization	TDM mode
T1 (Combo TDM OC48 FRU)	SAToP	Port (See Note 3.)	0	DS1 (See Note 2.)	ANSI
E1 (Combo TDM OC48 FRU)	SAToP	Port (See Note 3.)	0	E1	ETSI
T1 (TDM FRU)	CESoP	Port (See Note 3.)	1-24	DS0	ANSI
E1 (TDM FRU)	CESoP	Port (See Note 3.)	1-31	E0	ETSI

Note 1: When DS3 framing is set to M13, only the DS3->DS1 loopback is supported.
Note 2: If DS1/E1 channelization is configured, a PW must be set before enabling the remote loopback.
Note 3: For DS1 ports (also known as T1 ports), and for E1 ports, only port loopback is supported.
Note 4: For DS3 ports, port loopback is supported with DS1 channelization.
Note 5: For E3 ports, port loopback is supported.

This table shows the devices supported for TDM ports.

Table 10-2
TDM ports by device

Device	TDM SFP ports on the device	TDM ports on an FRU module (if equipped)
3903	2	Not applicable
3903x	2	Not applicable
3904	3	Not applicable
3905	3	Not applicable
3906	3	Not applicable
3924	7	Not applicable
3926	7	16 TDM ports on TDM FRU 14 TDM ports on Combo TDM OC48 FRU
3928	9	Not applicable
5142	23	Not applicable
5160	23	Not applicable

Copyright© 2022 Ciena® Corporation

TDM module ports

This table summarizes the TDM module ports supported in which mode, attachment circuit (AC) type and channelization.

Table 10-3
TDM module ports

TDM module port	Mode	TDM port	AC type	Channelization
DS1	ANSI	T1	CESoP	DS1, DS0
DS1	ANSI	T1	SAToP	DS1
E1	ETSI	E1	CESoP	E1, E0
E1	ETSI	E1	SAToP	E1

Combo TDM OC48 FRU module ports

This table summarizes the TDM module ports supported in which mode, attachment circuit (AC) type and channelization.

Table 10-4
Ports on the Combo TDM OC48 FRU module

Port number	Mode	Port type	AC type	Channelization
OC48-STM16-CEM (Port 3. See Note 2.)	ANSI	OC48, SFP	CESoP	STS48, STS12, STS3c, STS1
			SAToP	DS1
OC48-STM16-CEM (Port 3. See Note 2.)	ETSI	STM16, SFP	CESoP	VC4-16c, VC4
			SAToP	E1
OC12-STM4-CEM (Ports 3, 4, 5, and 6. See Note 2.)	ANSI	OC3/OC12, SFP	CESoP	STS12, STS3c, STS1
			SAToP	DS1
OC12-STM4-CEM (Ports 3, 4, 5, and 6. See Note 2.	ETSI	STM1/STM4, SFP	CESoP	VC4-4c, VC4
			SAToP	E1
OC3-STM1-CEM (Ports3, 4, 5, and 6. See Note 2.)	ANSI	OC3/OC12, SFP	CESoP	STS3c, STS1
			SAToP	DS1
OC3-STM1-CEM (Ports3, 4, 5, and 6. See Note 2.	ETSI	STM1/STM4, SFP	CESoP	VC4
			SAToP	E1

Copyright© 2022 Ciena® Corporation

Table 10-4
Ports on the Combo TDM OC48 FRU module

Port number	Mode	Port type	AC type	Channelization
DS3 (Ports 13-16.)	ANSI	DS3, DIN 1.0/2.3	CESoP	STS1
			SAToP	DS3, DS1
E3 (Ports 13-16.)	ETSI	E3, DIN 1.0/2.3	SAToP	E3
T1 (Ports 7-12.)	ANSI	DS1, RJ-45	SAToP	DS1
E1 (Ports 7-12.)	ETSI	E1, RJ-45	SAToP	E1

Note 1: Ports 1 and 2 are not used.

Note 2: If port 3 is used for an OC48 or STM16 connection, then ports 4, 5, and 6 cannot be used. If the FRU is in ANSI mode and if port 3 is not used for an OC48 connection, then ports 3, 4, 5, and 6 can each be used for an OC3 or OC12 connection, in any combination. If the FRU is in ETSI mode and if port 3 is not used for an STM16 connection, then ports 3, 4, 5, and 6 can each be used for an STM4 or STM1 connection, in any combination.

Port alarms

TDM SFP ports do not have this behavior and conform to the RFC 5086 and G.775 specifications.

Alarms on the ports of the Combo TDM OC48 FRU module are compliant to the following standards:

- GR-253-CORE (September 2000), section 6.2.1
- GR-820_CORE (December 1997)
- ANSI T1.231 (1997)

Signal degrade and signal fail

Signal degrade and signal fail is supported on OC3/STM1, OC12/STM4 and DS3 SFPs. A bit error rate from 3 to 9 is supported. The set threshold, clear threshold and interval cannot be directly set.

Copyright© 2022 Ciena® Corporation

This table lists the specific set of configuration values supported. The default is signal degrade and signal fail are not enabled.

Table 10-5
Signal degrade and signal fail values

Signal degrade or signal fail bit error rate value	Bit error rate clear threshold	Bit error rate set threshold	Interval
3 (10^{-3})	10	10	32
4 (10^{-4})	22	10	48
5 (10^{-5})	20	28	320
6 (10^{-6})	15	22	2400
7 (10^{-7})	15	22	24000
8 (10^{-8})	15	22	240000
9 (10^{-9})	15	22	2400000

Signal degrade and signal fail may be set on a TDM port using a new signal-degrade and signal-fail CLI attribute under the TDM port set command.

- When the signal degrade bit error rate set threshold is exceeded for the specified interval, a signal degrade fault is raised against the SFP.
- When the signal fail bit error rate set threshold is exceeded for the specified interval, a signal fail fault is raised against the SFP.
- Signal degrade and signal fail settings may be turned off by issuing an unset command.
- The port tdm port show command displays the signal degrade and signal fail bit error rate configured value when set.

Protection groups

Clock-mode is not specified on the TDM port for TDM 3926 ports. The clock-mode is internally set based on the clock selection algorithm of the protection group assigned to the TDM port. If there is no protection group, the clock mode defaults to internal.

Copyright© 2022 Ciena® Corporation

Encapsulation

This table summarizes encapsulation information for TDM services.

Table 10-6
Encapsulation information

Encapsulation type	Description	Supported on
Native MPLS	Native MPLS uses an MPLS header in front of the CEM encapsulation header -- SAToP RFC 4553, CEP RFC 4842, or CESoPSN RFC 5086 -- and the TDM payload. The bottom/innermost PW label provides the demultiplexing function. This encapsulation is compliant with the MPLS Label stack in IETF RFC 4553 and RFC 4842.	Supported on all TDM CEM SFPs, 3926 16xT1/E1 module ports, and 3926 3Gb TDM Combo module ports.
MEF8	This complies with Metro Ethernet Forum (MEF) specification MEF.8 for PDH service emulation. It uses an Ethernet header followed by the CEM encapsulation header -- SAToP RFC 4553 header, CEP RFC 4842 header, or CESoPSN RFC 5086 header – the optional RTP header, and the TDM payload. The CEM encapsulation header is also called MEF.8 Emulation header. The MEF.8 PDH CESoETH control word equals the SAToP or CESoPSN control word. The Ethernet header is: • Emulated Circuit Identifier (ECID) • EtherType 0x88d8 • VLAN (optionally a second VLAN) • MAC SA, DA. The DA is that of the far-end EC-ID endpoint.	Supported on TDM SFP ports, 3926 16xT1/E1 module ports, and 3926 TDM Combo module ports.
MEF8 over MPLS	The MEF.8 CES header (including MAC SA & DA) is encapsulated in an MPLS PW label, that is then transported over an MPLS LSP (with an LSP label) and optionally transported over an Ethernet VLAN, when transmitted over an Ethernet interface using a standard Ethernet header which includes another, outer MAC SA/DA. The use of the MPLS PW label is to transport a MEF.8 CES over an MPLS transport network. This PW label identifies the MEF.8 CES endpoint in the MPLS transport network. The inner MEF.8 CES header MAC DA is then used to switch the MEF.8 CES to the correct egress port for MEF.8 CES termination. The inner MEF.8 CES header EC-ID identifies the TDM Circuit ID being emulated. This is a non-standard technique that allows interop between a MEF.8 CES implementation and a Native MPLS TDM circuit emulation implementation.	Supported on all TDM CEM SFPs, and on 3926 16xT1/E1 Module ports.
Dry Martini	The encapsulation header is an MPLS header with an optional control word. There is no tunnel label.	Supported on 3926 T1/E1

Copyright© 2022 Ciena® Corporation

Transport

This table summarizes transport information for TDM services.

Table 10-7
Supported transport mechanisms

Transport mechanism	Header contains	PSN identified by means of	Supported on:
802.1Q	• service VLAN (S-VLAN)	• S-VLAN • destination MAC address	Combo TDM OC48 FRU module ports
Q-in-Q	• customer VLAN (C-VLAN) • S-VLAN	• S-VLAN • destination MAC address	TDM SFP ports
MPLS	tunnel label stacked on top of PWE label	• tunnel label	Combo TDM OC48 FRU module ports

E1 cable pin configuration

The following table lists the E1cable pin configuration.

Table 10-8
E1 cable pin configuration

X1		X2	
PIN	**Signal**	**Pin**	**Signal**
1	Rx -	4	Rx -
2	Rx +	5	Rx +
4	Tx -	1	Tx -
5	Tx +	2	Tx +

Copyright© 2022 Ciena® Corporation

E1 channel number mapping

This table describes the mapping of E1 channel numbers. This mapping allows users to create the right traffic from subtending equipment and test sets.

Table 10-9
E1 channel number mapping

IWF Channel Number	CSop STM-1 VC Mapping Scheme	ITU-T VC Mapping Scheme
0	1-1-1	1-1-1
1	2-1-1	2-1-1
2	3-1-1	3-1-1
3	1-2-1	1-2-1
4	2-2-1	2-2-1
5	3-2-1	3-2-1
6	1-3-1	1-3-1
7	2-3-1	2-3-1
8	3-3-1	3-3-1
9	1-4-1	1-4-1
10	2-4-1	2-4-1
11	3-4-1	3-4-1
12	1-5-1	1-5-1
13	2-5-1	2-5-1
14	3-5-1	3-5-1
15	1-6-1	1-6-1
16	2-6-1	2-6-1
17	3-6-1	3-6-1
18	1-7-1	1-7-1
19	2-7-1	2-7-1
20	3-7-1	3-7-1
21	1-1-2	1-1-2
22	2-1-2	2-1-2

Copyright© 2022 Ciena® Corporation

Table 10-9
E1 channel number mapping

IWF Channel Number	CSop STM-1 VC Mapping Scheme	ITU-T VC Mapping Scheme
23	3-1-2	3-1-2
24	1-2-2	1-2-2
25	2-2-2	2-2-2
26	3-2-2	3-2-2
27	1-3-2	1-3-2
28	2-3-2	2-3-2
29	3-3-2	3-3-2
30	1-4-2	1-4-2
31	2-4-2	2-4-2
32	3-4-2	3-4-2
33	1-5-2	1-5-2
34	2-5-2	2-5-2
35	3-5-2	3-5-2
36	1-6-2	1-6-2
37	2-6-2	2-6-2
38	3-6-2	3-6-2
39	1-7-2	1-7-2
40	2-7-2	2-7-2
41	3-7-2	3-7-2
42	1-1-3	1-1-3
43	2-1-3	2-1-3
44	3-1-3	3-1-3
45	1-2-3	1-2-3
46	2-2-3	2-2-3
47	3-2-3	3-2-3
48	1-3-3	1-3-3
49	2-3-3	2-3-3

Copyright© 2022 Ciena® Corporation

Table 10-9
E1 channel number mapping

IWF Channel Number	CSop STM-1 VC Mapping Scheme	ITU-T VC Mapping Scheme
50	3-3-3	3-3-3
51	1-4-3	1-4-3
52	2-4-3	2-4-3
53	3-4-3	3-4-3
54	1-5-3	1-5-3
55	2-5-3	2-5-3
56	3-5-3	3-5-3
57	1-6-3	1-6-3
58	2-6-3	2-6-3
59	3-6-3	3-6-3
60	1-7-3	1-7-3
61	2-7-3	2-7-3
62	3-7-3	3-7-3

T1 SFP pin mapping

The T1 interface terminates in an 8-pin RJ45 connector, wired in accordance with ANSI T1.403.

Table 10-10
T1 SFP (XCVR-TDM4GE) pin function

Pin	Signal
1	Tx-Ring
2	Tx-Tip
3 6 7 8	Ground
4	Rx-Ring
5	Rx-Tip

Copyright© 2022 Ciena® Corporation

TDM FRU pin function

This table describes the mapping of pinouts in an RJ45 connector.

Table 10-11
TDM FRU pin function

Pin	Signal
1	RTIP
2	R-Ring
4	TTIP
5	T-Ring
7 8	Ground
3	NC
6	NC

DS1 channel numbers

This table describes the mapping of DS1 channel numbers.

Note: The table shows the numbering of DS1s starting from 1, but Ciena's numbering starts from 0.

Table 10-12
DS1 channel number mapping

STM1Mux Port	STS-1	VT-Group	VT1.5	STM1Mux Port	STS-1	VT-Group	VT1.5
1	1	1	1	**43**	1	1	3
2	1	2	1	**44**	2	2	3
3	1	3	1	**45**	3	3	3
4	1	4	1	**46**	4	4	3
5	1	5	1	**47**	5	5	3
6	1	6	1	**48**	6	6	3
7	1	7	1	**49**	7	7	3
8	1	1	2	**50**	1	1	4
9	1	2	2	**51**	2	2	4
10	1	3	2	**52**	2	3	4

Copyright© 2022 Ciena® Corporation

Table 10-12
DS1 channel number mapping

STM1Mux Port	STS-1	VT-Group	VT1.5	STM1Mux Port	STS-1	VT-Group	VT1.5
11	1	4	2	**53**	2	4	4
12	1	5	2	**54**	2	5	4
13	1	6	2	**55**	2	6	4
14	1	7	2	**56**	2	7	4
15	1	1	3	**57**	3	1	1
16	1	2	3	**58**	3	2	1
17	1	3	3	**59**	3	3	1
18	1	4	3	**60**	3	4	1
19	1	5	3	**61**	3	5	1
20	1	6	3	**62**	3	6	1
21	1	7	3	**63**	3	7	1
22	1	1	4	**64**	3	1	2
23	1	2	4	**65**	3	2	2
24	1	3	4	**66**	3	3	2
25	1	4	4	**67**	3	4	2
26	1	5	4	**68**	3	5	2
27	1	6	4	**69**	3	6	2
28	1	7	4	**70**	3	7	2
29	2	1	1	**71**	3	1	3
30	2	2	1	**72**	3	2	3
31	2	3	1	**73**	3	3	3
32	2	4	1	**74**	3	4	3
33	2	5	1	**75**	3	5	3
34	2	6	1	**76**	3	6	3
35	2	7	1	**77**	3	7	3
36	2	1	2	**78**	3	1	4
37	2	2	2	**79**	3	2	4

Copyright© 2022 Ciena® Corporation

Table 10-12
DS1 channel number mapping

STM1Mux Port	STS-1	VT-Group	VT1.5	STM1Mux Port	STS-1	VT-Group	VT1.5
38	2	3	2	**80**	3	3	4
39	2	4	2	**81**	3	4	4
40	2	5	2	**82**	3	5	4
41	2	6	2	**83**	3	6	4
41	2	7	2	**84**	3	7	4

Copyright© 2022 Ciena® Corporation

Linear APS 1+1

Linear APS 1+1 protection (where APS stands for automatic protection switching) is used to protect the point-to-point multiplex section layer in a SONET/SDH network. It provides dedicated (non-shared) protection. The MSP protocol, at both ends of a multiplex section, requests and acknowledges switch action by using the APS bytes (K1 and K2 bytes in the MSOH of the protection section), as illustrated in this figure:

Figure 10-1
Linear APS 1+1 bidirectional protection switching

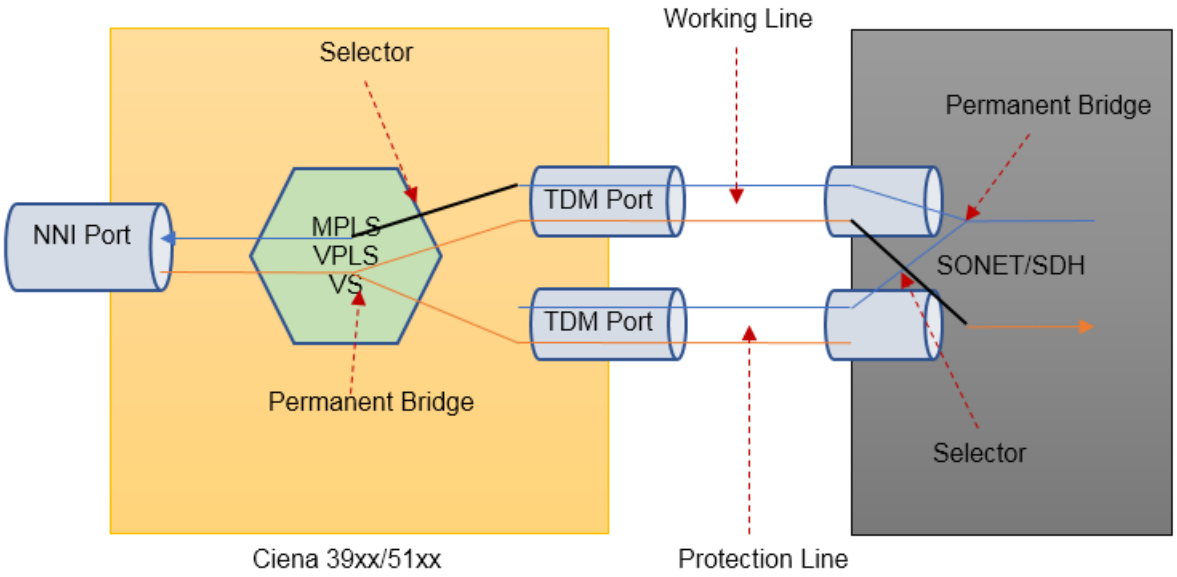

In linear APS 1+1 protection mode:

- At Tx direction, the working line is permanently bridged to the protection line.
- At Rx direction, the selector selects either the working line or the protection line (that is, selectively blocks either the protection line or the working line).

Copyright© 2022 Ciena® Corporation

The devices that support the linear APS 1+1 feature and the possible number of linear APS 1+1 protection groups that can be configured (assuming each protection group consumes two TDM SFP ports) are summarized in this table:

Table 10-13
Linear APS 1+1 devices supported

Device	TDM SFP ports	Linear APS 1+1 support	Possible number of linear APS 1+1 protection groups*	Notes
3903	2	N/A	N/A	No MPLS support
3903x	2	N/A	N/A	No MPLS support
3904	3	N/A	N/A	No MPLS support
3905	3	N/A	N/A	No MPLS support
3906	3	N/A	N/A	
3924	7	Yes	10	
3926	7	Yes	10	
3928	9	Yes	10	
5142	23	Yes	10	
5160	23	Yes	10	
* This number is constrained on each device by the number of physical ports available and one port reserved as the NNI port.				

Linear APS 1+1 protection supports the same scale number of ANSI/ETSI PW as is supported on various 6.x platforms (see "TDM ports by device" on page 10-5 for details).

APS 1+1 on the Combo TDM OC48 FRU module

If the Combo TDM OC48 FRU module is operating in ANSI tdm mode (the default tdm mode), then the ports on the FRU can support APS 1+1 for port protection. On the FRU, the user can configure OC12 and OC3 ports on ports 3, 4, 5, and 6. (If the Combo TDM OC48 FRU module is set to ETSI tdm mode, APS 1+1 protection is not supported on its ports.)

The following figure illustrates the high-level operation of linear APS 1+1 protection switching on the Combo TDM OC48 FRU module with native MPLS SFPs.

Copyright© 2022 Ciena® Corporation

Figure 10-2
Linear APS 1+1 on the Combo TDM OC48 FRU module

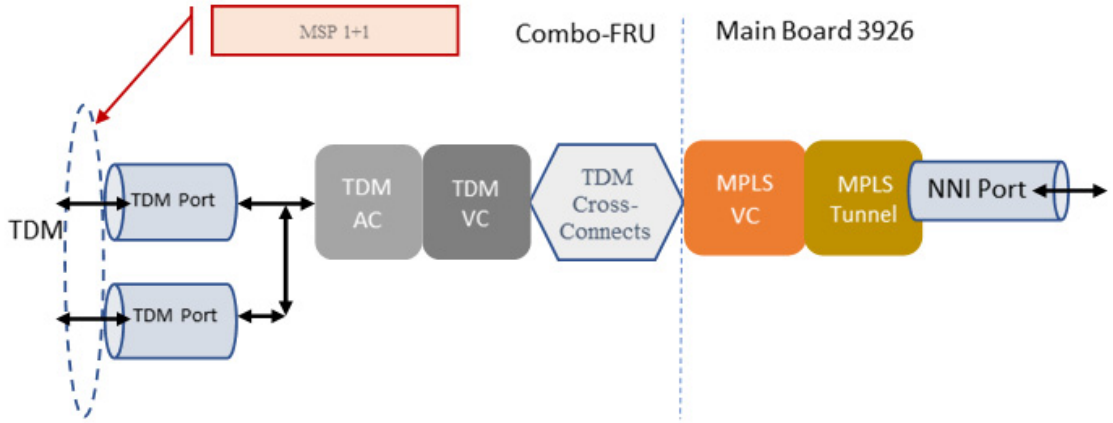

In SAOS 6.19 the APS 1+1 working and protection ports must both be on the Combo TDM OC48 FRU module.

APS 1+1 can be supported on OC3/OC12 ports with CESoP as well as SAToP attachment circuits (irrespective of the channelization).

Table 10-14

TDM Module Port	TDM port	AC type	Channelization	Mode	Encapsulation	APS 1+1
OC3 port	OC3	CESoP	STS1, STS3c	ANSI	Native MPLS	Supported
		SAToP	DS1	ANSI	Native MPLS	Supported
OC12 port	OC12	CESoP	STS1, STS12c, STS3c	ANSI	Native MPLS	Supported
		SAToP	DS1	ANSI	Native MPLS	Supported
On the Combo TDM OC48 FRU module, if port 3 is not configured as an OC48 port, then the user can configure OC3 and OC12 ports on ports 3, 4, 5, and 6 of the Combo TDM OC48 FRU module. OC3 and OC12 ports can be configured in any combination.						

APS 1+1 can be configured on Combo TDM OC48 FRU ports 3, 4, 5, and 6 with native TDM SFPs for protecting the TDM traffic with native MPLS encapsulation.

Copyright© 2022 Ciena® Corporation

Table 10-15
Combo TDM OC48 FRU module MPLS encapsulation support

Encapsulation	MPLS LSP	MPLS PW	SFP type	APS 1+1	Release
Native MPLS	Static and dynamic	Static only	Native only	Supported	6.19.0
MEF.8	Not applicable	Not applicable	Not applicable		
Dry Martini	Not applicable	Not applicable	Not applicable		

APS 1+1 feature restrictions

Creation, deletion and modification of linear APS protection groups is supported with the following restrictions:

- The protection port must be the same type and have the same configuration as the working port for the following attributes:
  - mac
  - clock-mode
  - c-tpid
- If APS 1+1 is configured on a CSoP SFP, the protection port can have different/independent configuration from the working port for the following attributes:
  - section-trace
  - signal-degrade
  - signal-fail
  - loopback
  - prbs

  Note: If the Combo TDM OC48 FRU module is operating in ANSI mode (the default tdm mode), and if APS 1+1 is configured on its ports, loopback and prbs cannot have independent configurations on the protection port.
- Once a port is added as a protection port in a protection group, the following port configuration cannot be changed until the port is removed from the protection group:
  - mac
  - clock-mode
  - c-tpid
- When the following port configuration of a working port is changed, it is automatically applied to the protection port in the same protection group:
  - mac
  - clock-mode
  - c-tpid

Copyright© 2022 Ciena® Corporation

- When a protection port is added to a protection group, it must not have any TDM attachment circuit (AC) configuration associated with the port.
- When creating a TDM AC and associating it with a port, if the port is a member of a protection group, the port must be the working port of that protection group.
- The protection port can be removed from a protection group even if there is already a TDM AC associated with the corresponding working port in the same protection group.
- The working port can be removed from a protection group even if there is already a TDM AC associated with the port.
- A TDM port cannot be deleted if it is a member of a protection group.
- A protection group port membership can only be changed if it is administratively disabled.
- A protection group can only be enabled if both working and protection ports are added.
- A protection group can only be deleted after all the member ports are removed
- A protection group can only be deleted after it is administratively disabled
- A protection group must be disabled before the changing the “revertive” attribute.
- The following attributes of a protection group can be changed (on the fly) when the protection group is enabled:
  - wait-to-restore
  - hold-off
- APS 1+1 protection is not supported on the Combo TDM OC48 FRU if the tdm mode of the FRU is set to ETSI.

In addition to configuration information, the linear APS protection group provides the following runtime statistics:

- Total number of Far-End Protection-Line Failures (FEPLF)
- Total number of remote timeouts (unacknowledged requests leading to timeout)

Linear APS modes of operation

Different modes of operation are used to describe linear APS. This feature supports at minimum the 1+1 architecture, bidirectional mode and revertive switching. (For detail of the linear MSP APS protocol operations, refer to ITU-T G.841.)

Copyright© 2022 Ciena® Corporation

1+1 architecture

Two main linear APS architectures are defined by governing standards: the 1+1 architecture and 1:N architecture. The 1+1 architecture has one protection line for one working line. In addition, the signal is permanentlybridged from the head-end onto both working and protection lines in normal working conditions. At the tail end, both the working and protection lines are monitored independently and one of the lines is selected based on signal integrity, local switch commands, and far-end switch requests.

The 1:N architecture has one protection line for N working lines. A special case of the 1:N architecture is 1:1, in which there is one protection line for one working line. In the case of 1:N, however, the data is transmitted only on the working line.

For this feature, only the 1+1 architecture is supported. 1:N and 1:1 architectures are not supported. The linear APS protection group defaults to the 1+1 architecture—the architecture is not provisionable given other architectures are not supported.

Protection mode

A system using the 1+1 architecture operates, by default, in a unidirectional mode. In this mode of operation, the switching is complete when a line in the failed direction is switched to the protection line. In bidirectional switching mode, a line is switched to the protection line in both directions. This is accomplished through the use of APS signaling.

For this feature, at minimum, bidirectional mode is supported. To comply with standards, unidirectional protection mode is also supported when operating with a remote device that is configured for unidirectional operation. Provisioning support is not available to change the mode from bidirectional to unidirectional and vice versa in the initial release.

To support bidirectional switching, the equipment at both ends must be provisioned to operate in bidirectional mode. In the case of mismatch detected through the K2 byte signaling, the operational mode reverts to unidirectional switching (this implies that unidirectional mode needs to be supported in order to be standard compliant). In unidirectional mode, APS signaling is not required for switching but is required to determine the mode of operation and local switch action.

In bidirectional mode, only the 2-phase protocol operation (refer to ITU-T G.808.1 Clause 24 for more detail of APS protocol phases) is supported. There is no support for 1-phase and 3-phase protocol operations (per ITU-T G.808.1) in the initial release. However, the implementation is compatible with remote equipment operating in 1-phase or 3-phase protocols.

Copyright© 2022 Ciena® Corporation

Revertive switch mode

In non-revertive mode, traffic remains on the protection line until another switch request is received. In revertive mode, when the condition that caused a switch to the protection line has been cleared, traffic is switched back to the working line after the WTR time has expired. In a 1+1 system, the default mode is non-revertive switching.

For this feature, at minimum, revertive switching is supported. Both revertive and non-revertive modes are provisionable and supported since the default mode of operation is non-revertive switching, as mentioned previously.

Linear APS switchover criteria

The switch initiation time is the time the equipment takes to detect a signal fail (SF) or signal degrade (SD) condition and to initiate a protection switch. For SF conditions, the switch initiation time is within 10 ms.

The time to complete a switch once initiated is within 50 ms. If both local and remote devices detect LOS at the same time, switchover time can be longer.

> ***Note:*** The switch completion time excludes the switch initiation time mentioned previously.

Copyright© 2022 Ciena® Corporation

Linear APS switch commands

Available linear APS switch commands are:

Table 10-16
APS switch commands

Command	Description
Lockout of Protection	This command prevents the working line from switching to the protection line (the command is only available on the protection line). This feature may support the lockout of protection command in a future release and the lockout of protection will persist across node reboot. Minimally, this feature supports a remote-initiated lockout of protection in bidirectional mode.
Forced Switch	This command is used to force a protection switch: • From the working line to the protection line and • From the protection line back to the working line This feature supports the forced switch command and the forced switch persists across node reboot.
Manual Switch	This command is used to perform a manual protection switch: • From the working line to the protection line and • From the protection line back to the working line This feature may support the manual switch command in a future release and the manual switch will persist across node reboot. Minimally, this feature supports a remote-initiated manual switch in bidirectional mode.

Copyright© 2022 Ciena® Corporation

Table 10-16
APS switch commands

Command	Description
Clear	This command is used to clear a user-initiated local switch command on a specific line. This feature supports the clear command. (There is no need to persist the clear command across node reboot by definition). ***Note:*** As stated in G.841 section 7.1.2.1, this command does not clear the WTR timer following an automatic switch in revertive mode. To restore traffic to the working line following an automatic switch in non-revertive mode or before the WTR timer has expired in revertive mode, the forced switch command may be issued followed by the clear command.
Wait to Restore (WTR)	After the condition which triggered the protection switch has cleared, the wait to restore period (in seconds) must elapse before traffic is reverted. This is intended to avoid rapid switch oscillations. GR-253-CORE specifies a Wait to Restore period of 5 to 12 minutes. G.783 defines the same range but also allows for a short WTR period. This feature supports provisionable WTR time in the previously mentioned range with a default WTR time of 5 minutes.
Exercise	This command exercises the protection line for a protection switch of the specified working line by issuing the exercise request for that working line and checking the response on the APS channel. This feature may support the exercise command in a future release. (There is no need to persist the exercise command across node reboot by definition). Minimally, this feature supports a remote-initiated exercise command in bidirectional mode.

Copyright© 2022 Ciena® Corporation

APS status

APS status and associated statistics counters are used to indicate when the following conditions are detected:

Table 10-17
APS status

Condition	Description
APS Mode Mismatch	K2 byte bit 5 indicates the protection architecture and K2 bits 6-8 indicate the mode (unidirectional or bidirectional). A mode mismatch condition occurs when there is a conflict between the current local mode and the received K2 mode information. This feature support displaying the APS mode mismatch status and its associated statistics counter.
APS Channel Mismatch	This condition occurs when there is a mismatch between the transmitted K1 byte channel and the received K2 byte channel. This feature supports displaying the APS channel mismatch status and its associated statistics counter.
Protection Switch Byte Failure (PSBF)	This condition occurs when either an inconsistent APS byte or an invalid code is detected. An inconsistent APS byte occurs when no three consecutive K1 bytes of the last 12 successive frames are identical. An invalid code occurs when the incoming K1 byte contains an unused code or a code irrelevant for the specific switching operation. This feature supports displaying the protection switch byte failure status and its associated statistics counter.
Far-End Protection-Line Failure (FEPLF)	This condition is declared based on receiving SF on the protection line in the K1 byte. This feature supports displaying far-end protection-line failure status and its associated statistics counter.

APS protocol signaling and K bytes

The APS protocol signaling is carried within the K1 and K2 bytes of the MSOH (Multiplex Section Overhead) and the line overhead (SONET). The two ends of the Linear APS use these bytes to exchange requests and acknowledgments for protection switch actions. Other APS information such as APS architecture and mode of operation is also contained within these two bytes. The K bytes are generated on the protection line. (Refer to ITU-T G.841 sub-clause 7.1.1 for details of how K bytes are used.)

Copyright© 2022 Ciena® Corporation

K1 byte semantics

The K1 byte is used to indicate a request by a channel for a switch action. The K1 byte is divided into two parts. Bits 1 through 4 are used to indicate the request, while bits 5 to 8 are used to indicate the channel making the request.

The APS request is encoded with the first four bits of the K1 byte and are defined with a relative priority. Only one APS request (highest priority) is expected to be generated at any given time. The APS request can be further divided to indicate the following: a condition such as signal failure (SF) or signal degrade (SD), a state such as "wait-to-restore" and "do not revert", and external requests such as forced and manual switch from an operator.

The following table indicates the local conditions that are supported by this feature. However, it is expected that the local node responds to all defined requests coming from the far-end to support compatibility with other vendors running the APS protocol. For SF (signal failure) and SD (signal degrade) conditions detected locally, the APS-generated request uses low priority encoding (in accordance to GR-253 -CORE).

Bits 1 through 4 of the K1 byte are used to indicate the request, as listed in this table:

Table 10-18
K1 byte request field

Bit 1234	Switch and state request
1111	Lockout of Protection
1110	Forced Switch
1101	SF (High Priority)
1100	SF (Low Priority)
1011	SD (High Priority)
1010	SD (Low Priority)
1001	(unused)
1000	Manual Switch
0111	(unused)
0110	Wait-To-Restore
0101	(unused)
0100	Exercise
0011	(unused)

Copyright© 2022 Ciena® Corporation

Table 10-18
K1 byte request field

Bit 1234	Switch and state request
0010	Reverse Request
0001	Do Not Revert
0000	No Request

Bits 5 through 8 of the K1 byte are used to indicate the number of the channel that is making the request, using the following encoding:

Table 10-19
K1 byte channel field

Bit 5678	Channel and notes
0000	Null Channel • SD and SF requests apply to conditions detected on the protection line. • For 1+1 systems, forced and manual switch requests apply to the protection line. • Lockout of protection requests.
0001 to 1110	Working Channel • Codes 1 through n apply in 1:n architecture. • For 1+1 architecture, only code 0001 applies.
1111	Extra Traffic Channel • Only applies to 1:n architecture. • Also used with “No Request”.

Note: For 1+1 architecture, the working channel is always set to 0001.

Copyright© 2022 Ciena® Corporation

K1 byte generation

Local SF and SD conditions, wait-to-restore or do-not-revert states, and the external request are evaluated by a priority logic based on the descending order of request priorities shown in the previous table. A higher priority request replaces the current local request.

Table 10-20
K1 byte generation in operations

Operation	Description
Bidirectional	The priority of the local request is compared with the remote request on the received K1 byte, based on the descending order of request priorities. The sent K1 byte indicates: • Reverse request – if the remote request has a higher priority – if the requests are the same priority and the sent K1 byte already indicates a reverse request • Local request in all other cases
Unidirectional	The local request is always sent with the K1 byte.
In revertive mode	In the revertive mode of operation, when the condition (SF or SD) that originally triggered the protection switch to the protection line has cleared, a local wait-to-restore state is activated that triggers the procedure to revert/switch traffic back to the working line. In the absence of higher priority requests (such as forced switch), the wait-to-restore request is the highest priority request and is indicated in the transmitted K1 byte. Traffic is maintained on the protection line until the wait-to-restore timer expires, at which point traffic is restored to the working line. At this point, a no-request/null channel request is sent out with the K1 byte to indicate the switch back to the working line. In the non-revertive mode of operation, when the working channel is no longer in SF or SD condition, traffic is maintained on the protection line by the transmit of the do-not-revert request in the K1 byte. Both the wait-to-restore and do-not-revert requests are acknowledged by a reverse request in the K1 byte response. A no request in the K1 byte is acknowledged with a no-request response in the received K1 byte.

K2 byte semantics

Bits 1 to 4 are used to indicate the status of the bridge in the MSP switch. These bits are encoded in the same way as Bits 5 to 8 of the K1 byte. In 1+1 APS, the working channel is permanently bridged to both the working and protection lines. The K2 channel bits in this case are used as part of the acknowledgment of a switch request.

Copyright© 2022 Ciena® Corporation

Bit 5 of the K2 byte is used to indicate the architecture (1+1 or 1:1) as provisioned with the encoding of “0” indicating the supported architecture is 1+1, and the encoding of “1” indicating a 1:n architecture. In the event that a received mode as indicated in this bit does not match what is locally provisioned, an “APS Mode Mismatch" notice is raised against the protection group. It is expected that the side not operating as “1+1” modifies its operation to interwork as “1+1”.

Bits 6 to 8 of the K2 byte are used to indicate the provisioned protection mode with the encoding of “101” for bidirectional mode and “100” for unidirectional mode. The encoding of “111” is used to detect the incoming alarm indication signal line (AIS-L) in SONET and “110” used to detect the recirculating delay line (RDL-L). It is expected that the TDM port raises the AIS-L and RDL-L conditions.

This feature supports displaying the most recently transmitted and received K bytes for each protection group.

K2 byte generation

The sent K2 byte indicates, in bits 1 to 4, for all architectures and operation modes, a null signal (0) if the received K1 byte indicates null signal and the extra traffic is not bridged. For all other cases, the number of the signal that is bridged is indicated by the sent K2 byte.

The sent K2 byte indicates in bit 5:

- 0 if 1+1 architecture
- 1 if 1:n architecture.

Bits 6 to 8 of the K2 byte are used to indicate the provisioned protection mode with the encoding of “101” for bidirectional mode and “100” for unidirectional mode.

Linear APS 1+1 caveats and limitations

The caveats and limitations for linear APS 1+1 APS include:

- SNCP 1+1 (sub-network protection) is not supported.
- Linear APS 1+1 dual-homing/dual-node is not supported.
- MSP 1:N and 1:1 architectures are not supported.
- Linear APS 1+1 can only be supported on MPLS VPLS VS with TDM smart SFPs.
- Only TDM over MEF8 over MPLS PW encapsulation is supported with TDM smart SFPs.

Copyright© 2022 Ciena® Corporation

- One and only one linear APS 1+1 protected TDM VC can be attached to the MPLS VPLS VS (this is limited by the implementation where it is relying on MPLS VPLS's L2 flooding mechanism to bifurcate traffic from NNI to UNI).

Caveats and limitations applying to APS 1+1 on the Combo TDM OC48 FRU module

The caveats and limitations applying to linear APS 1+1 on the Combo TDM OC48 FRU module are as follows:

- If APS 1+1 is configured on the Combo TDM OC48 FRU module, the working and protection ports must both be on the FRU module.
- The feature is supported only with "native MPLS" encapsulation.
- The feature is supported only with static MPLS pseudo wires with static and dynamic LSPs.

Framing

Some TDM ports allow the user to specify the framing type. Framing type is specific to the type of TDM port and is not always supported.

For DS1 the following framing types are supported:

- Unframed
- Super-Frame
- Extended-Super-Frame

For E1 the following framing types are supported:

- Unframed
- Basic
- Crc-Multiframe

For DS3 the following framing types are supported:

- Unframed
- C-Bit
- M13

For E3 the following framing types are supported:

- Unframed
- G.751

Channel Associated Signaling (CAS) for nxDSO service

When in nxDS0 mode, Channel Associated Signaling (CAS) is a method of signaling each nxDS0 traffic channel rather than having a dedicated signaling channel (like ISDN). CAS is also referred to as Robbed Bit Signaling. In this

Copyright© 2022 Ciena® Corporation

type of signaling, the least significant bit of information is “robbed” from the channels that carry voice and is used to transmit framing and clocking information. This is sometimes called “in-band” signaling.

nxDS0 mode is currently supported only on the 16xT1/E1 FRU, and this FRU supports a CAS-transport configurable mode.

nxDS0 CAS-transport configuration notes

The following are configuration notes for the TDM nxDS0 CAS-transport implementation on the 3926 16xT1/E1 FRU.

When a 3926 16xT1/E1 FRU port is in the nxDS0 channelized mode:

- CAS-transport mode auto-configures the Packet Delay Variation (PDV) when CAS is enabled, regardless of the current PDV value or if a PDV is specified. The new PDV is saved as part of the configuration.
- The PDV is not restored to the previous value when CAS-transport is disabled.
- The PDV may be changed using the set command after CAS-transport is enabled, but a warning is generated indicating that this is not recommended because it may cause signaling issues between the DTE and the DCE.
- When CAS-transport is enabled, payload size in the TDM profile must be set using the following formula: DS1:24 * N E1:16 * N (where N is the number of timeslots). This restriction is because the packetization latency must be 3 ms for DS1 and 2 ms for E1 and this is the only payload size that results in that latency.

  Note: The actual formula includes “+ (N+1)/2”, but this is unnecessary because the simpler formula already results in the correct packetization latency.
- CAS-transport can be enabled only on ports that do not use a framing type of unframed.
- When CAS-transport and non-CAS-transport VCs are simultaneously in use on the same T1/E1, PRBS is only applied to the CAS-transport portion (56KHz) of the signal for all the VCs. This means that for non-CAS-transport VCs (64KHz), a portion of that signal is not filled w/ PRBS signal.

Feature limitations

The following sections list feature limitations.

In addition to the limitations listed in the following sections, there are limitations on the features supported by TSOP transceivers, which can be used on SFP ports. For details, see Table 10-22.

Copyright© 2022 Ciena® Corporation

Ethernet feature restrictions

These features are restricted and are not configurable on a port that has an SFP TDM port associated with it:

- Benchmark
- Aggregation
- LLDP
- EOAM
- Dot1X
- RSTP/MSTP
- VLLI
- CFM MEP
- Sync E

These features are not supported on a port that has an SFP TDM associated with it. When the TDM port associated is deleted, the physical port assumes the attributes of these features and are defined for the physical port.

- TWAMP
- virtual switch
- sub-ports
- Private Forwarding Groups
- traffic profiling
- egress queuing
- broadcast containment
- VLAN

Section trace

Section trace is only available on optical devices.

Pseudo-Random Bit Sequence

Pseudo-Random Bit Sequence (PRBS) insertion/monitoring is available on VCOP, CSOP and TPOP SFPs, T1/E1 3926 ports, and Combo TDM OC48 FRU ports.

It is not necessary to create a pseudowire to permit PRBS insertion and monitoring of a line channel.

On DS3, PRBS insertion and monitoring is only available on the network side.

If the PRBS direction is not the same between the insertion and monitoring, the channeling must be the same.

Copyright© 2022 Ciena® Corporation

Framed bit error-rate testing (BERT) and framed SAToP services cannot work simultaneously. This impacts PBRS as follows:

- PBRS monitoring is not supported for DS3-E3 ports supporting SAToP services if framing is configured. (DS3-E3 ports supporting SAToP services are supported on the Combo TDM OC48 FRU module.)
- PRBS is not supported for DS3 ports supporting SAPoP services if the attachment-circuit channelization is DS1. (DS3 ports supporting SAPoP services are supported on the Combo TDM OC48 FRU module.)
- PRBS line monitoring is not supported for a T1 ports that is supporting SAToP services if the framing for the TI port is set to “super-frame” or to “extended-super-frame”.
- PRBS is supported for a DS3 or E3 port supporting SAToP services if the attachment-circuit channelization type is identical to the port type (DS3 channelization for a DS3 port or E3 channelization for an E3 port), and if the framing for the port is set to “unframed”.
- PRBS is not supported for E1 ports with MEF8 SAToP services configured if framing is PCM30 or PCM30-CRC. (PRBS is supported for E1 ports with MEF8 SAToP services configured if framing is PCM31 or PCM31-CRC.)

The following limitations apply to PRBS on the ports of the Combo TDM OC48 FRU module:

- On the Combo TDM OC48 FRU module, TDM ports support PRBS only towards the port. PRBS is currently not supported on the network side.
- The Combo TDM OC48 FRU does not support PRBS for DS3->STS1 in either direction.

Double tags

The TPID used for CVLAN tag is fixed at 0x8100, except for T1/E1 3926 ports which can be set at 0x8100, 0x9100 or 0x88a8.

If a C-VID is used on one virtual circuit, then a C-VID must be used on all the virtual circuits in the SFP.

There is no tag validation on single or double tags.

MPLS encapsulation

A list of restrictions is included in “MPLS encapsulation of MEF8 TDM PWs” on page 7-61. The restrictions differ while using a TDM SFP port.

Copyright© 2022 Ciena® Corporation

Timing

For any TDM SFP interworking with the TDM module port, the TDM module port must have the RTP header on, and the clocking mode must be internal on both ports. The 3926 supports all timing modes, TDM ports and SFP port configuration except for DCR. The 3903, 3903x, 3906, and 3924-901 (non-sync variant) platforms supports only ACR clocking mode.

Packet delay variation

The packet delay variation (PDV) value is not user-configurable for the ports on the Combo TDM OC48 FRU Module.

Copyright© 2022 Ciena® Corporation

Transceiver-related limitations

From SAOS 6.20 onwards, transparent SONET/SDH over packet (TSOP) transceivers can be used in the 3924, 3926, 3928, 5142, and 5160 platforms.

The TSOP transceivers are as follows:

- XCVR-TSOP03. Supports OC48 or STM16.
- XCVR-TSOP02. Supports OC12 or STM4.
- XCVR-TSOP01. Supports OC3 or STM1.

For instructions for creating and enabling the OCn or STMn ports to be used by TSOP transceivers, see Procedure 10-6, “Creating and enabling a port to be used by a TSOP transceiver”.

The following limitations apply to TSOP transceivers.

- Interoperation of a TSOP SFP with a non-TSOP SFP or port is not supported.
- Only port non-transparent terminal and facility loopback are supported.
- Direct back-to-back connecting TSoPs to each other optically is not supported.
- TSOP transceivers cannot be used on TDM ports on the Combo TDM OC48 FRU.
- If a TSOP transceiver is plugged in to a port, the system blocks certain CLI commands and certain command options because the transceiver is incompatible with the functionality. The commands and options are listed in Table 10-21.

Copyright© 2022 Ciena® Corporation

Table 10-21
Commands and options blocked for TSOP transceivers

Command	What is blocked
port tdm oc3-stm1 prbs set	Command
port tdm oc3-stm1-cem prbs set	Command
port tdm oc3-stm1 prbs unset	Command
port tdm oc3-stm1-cem prbs unset	Command
port tdm oc12-stm4 prbs set	Command
port tdm oc12-stm4-cem prbs set	Command
port tdm oc12-stm4 prbs unset	Command
port tdm oc12-stm4-cem prbs unset	Command
port tdm oc48-stm16-cem prbs set	Command
port tdm oc48-stm16-cem prbs unset	Command
port tdm oc3-stm1 create	port-mac option
port tdm oc12-stm4 create	port-mac option
port tdm oc48-stm16 create	port-mac option
virtual-circuit tdm-mef8 create	Command
virtual-circuit tdm-mef8 set	Command
virtual-circuit tdm-mef8 unset	Command
virtual-circuit tdm-mef8 delete	Command
virtual-circuit tdm-mef8 show	Command
port tdm oc3-stm1 set	c-tpid option
port tdm oc3-stm1 unset	c-tpid option
port tdm oc12-stm4 set	c-tpid option
port tdm oc12-stm4 unset	c-tpid option

Copyright© 2022 Ciena® Corporation

The following table summarizes the features supported by the TSOP transceivers.

Table 10-22
Feature supported and not supported by TSOP transceivers

Feature	XCVR-TSOP01 (OC3 / STM1)	XCVR-TSOP02 (OC12 / STM4)	XCVR-TSOP03 (OC48 / STM16)
RTP header	Always on (See Note 1.)	Always on (See Note 1.)	Always on (See Note 1.)
Clocking mode	DCR	DCR	DCR
Mode	ANSI / ETSI	ANSI / ETSI	ANSI / ETSI
AC type	SAToP	SAToP	SAToP
Channelization	STS3c	STS12c	STS48c
Packet delay variation (See Note 3.)	2.6 ms (default)	667 us (default)	1.3 ms (default)
Fixed payload-size	810 bytes	810 bytes	810 bytes
Packetization delay	41.7 ms	10.4 us	2.6 ms
MPLS	Supported	Supported	Supported
MEF8	Not supported	Not supported	Not supported
PRBS	Not supported	Not supported	Not supported
Port Mac	Not supported	Not supported	Not supported
C-TPID	Not supported	Not supported	Not supported
Signal-degrade	Not supported	Not supported	Not supported
Signal-fail	Not supported	Not supported	Not supported
SSM	Not supported	Not supported	Not supported
Loopcode	Not supported	Not supported	Not supported
CAS	Not supported	Not supported	Not supported

Note 1: The RTP header should always be on because the clocking mode is DCR (only), so RTP is required.

Note 2: In SAOS 6.21, the RTP header should always be off because the clock mode is ACR by default for 3903, 3903x, 3906, and 3924-901(non-sync variant) platforms.

Note 3: In SAOS 6.20, packet-delay variation is not configurable, and will use the defaults.

Copyright© 2022 Ciena® Corporation

Procedures

The procedures in this section are inter-related and are intended to be used as building blocks that can be combined to offer various TDM services over PSN. Note that there are procedures that are specific to one port type and procedures that are required for both port types. Refer to the flowcharts.

This figure shows the flow of procedures for end-to-end configuration of TDM SFP ports.

Copyright© 2022 Ciena® Corporation

Figure 10-3
TDM SFP port configuration

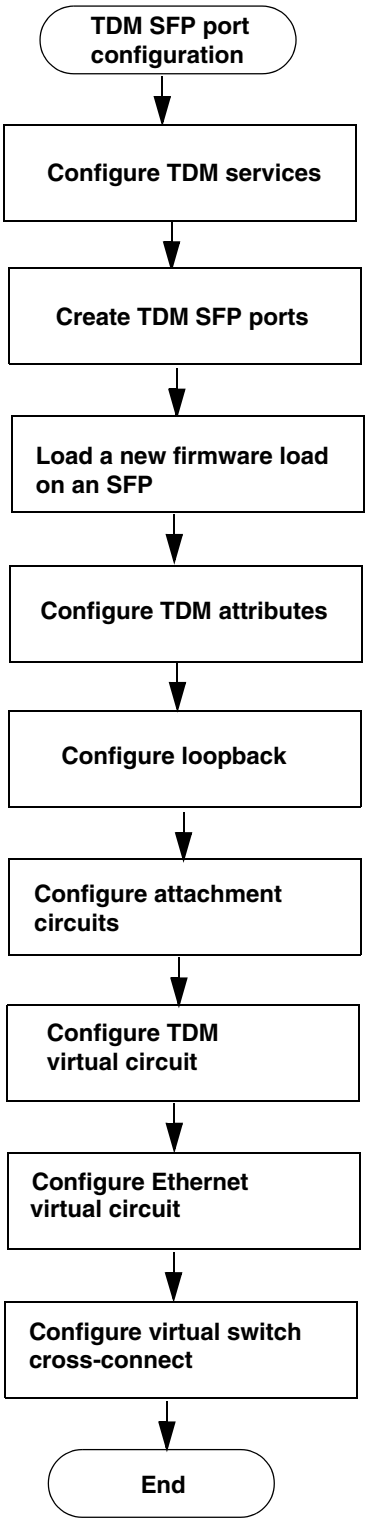

Copyright© 2022 Ciena® Corporation

Procedures for configuring TDM SFP ports are:

- “Configuring the FRU type” on page 10-43
- “Configuring TDM services” on page 10-45
- “Creating TDM SFP ports” on page 10-48
- “Deleting TDM SFP ports” on page 10-54
- “Loading a new firmware load on an SFP” on page 10-57
- “Creating and enabling a port to be used by a TSOP transceiver” on page 10-61
- Configure attributes:
  - “Changing the section trace” on page 10-63
  - “Configuring a protection group” on page 10-64
  - “Configuring a line-code on a T1 module” on page 10-65
  - “Configuring a line-code on an E1 module” on page 10-66
  - “Configuring framing on a T1 module” on page 10-67
  - “Configuring framing on a E1 module” on page 10-68
  - “Configuring framing on a Combo TDM OC48 FRU module” on page 10-69
  - “Configuring SSM” on page 10-71
  - “Configuring loopcode transmission on T1 module ports” on page 10-73
  - “Configuring C-TPID on the 3926 T1/E1 port” on page 10-75
  - “Configuring C-TPID on the TDM SFP” on page 10-76
  - “Configuring the cable length” on page 10-77
  - “Configuring the line build out cable length” on page 10-78
  - “Configuring TDM SFP PRBS insertion and monitoring” on page 10-79
- Configure signal degrade and signal fail
  - “Configuring signal degrade” on page 10-104
  - “Configuring signal fail” on page 10-105
- Configure loopback:
  - “Configuring OC3 port loopbacks” on page 10-106
  - “Configuring OC12 port loopbacks” on page 10-109
  - “Configuring OC48 port loopbacks” on page 10-112
  - “Configuring E1 port loopbacks” on page 10-115
  - “Configuring T1 port loopbacks” on page 10-116

Copyright© 2022 Ciena® Corporation

Global profiling enhancements

These topics were updated:

- “Global traffic profiles” on page 8-11
- “Recovering the clock” on page 10-85
- “Configuring TDM global synchronization” on page 10-91
- “Configuring TDM synchronization frequency protection groups” on page 10-92
- “Configuring TDM profiles” on page 10-96
- “Configuring attachment circuits” on page 10-100
- “Configuring TDM virtual circuits” on page 10-118
- “Configuring Ethernet virtual circuits” on page 10-137
- “Configuring virtual switch cross-connections” on page 10-138
- “Configuring jitter buffer monitoring on a T1 module” on page 10-139

You can display information about configured TDM services. Procedures to display information are:

- “Displaying information about the FRU module in 3926” on page 10-142
- “Displaying FPGA information about FRU modules” on page 10-143
- “Displaying TDM port information” on page 10-145
- “Displaying attachment circuit information” on page 10-152
- “Displaying TDM port loopbacks” on page 10-153
- “Displaying virtual circuit information” on page 10-156
- “Displaying virtual switch information” on page 10-158
- “Displaying virtual switch cross-connection information” on page 10-159

Additionally, this section provides end-to-end procedures for configuring Mobile Wireless Backhaul (MWB) in Metro Ethernet or MPLS networks:

- “Configuring SAToP services over 802.1Q Metro Ethernet on a TDM SFP or FRU port” on page 10-160
- “Configuring CESoP services over 802.1Q Metro Ethernet on a TDM SFP or FRU port” on page 10-162
- “Configuring SAToP services over Q-in-Q Metro Ethernet on a a TDM SFP or FRU port” on page 10-164
- “Configuring SAToP pseudowire over an MPLS network” on page 10-167

Copyright© 2022 Ciena® Corporation

Procedure 10-1
Configuring the FRU type

If you intend to remove a field replaceable unit (FRU) and replace it with an FRU of the same type, it is not necessary to remove power from the platform. However, it may be necessary to remove power from the FRU module before you remove it, as stated in the caution box following this paragraph. Once the second FRU module is inserted and powered up, it comes up with the same configuration as the original module and continues to pass traffic.

CAUTION
Risk of damage to equipment
Do not remove a Combo TDM OC48 FRU module from a 3936 unless the FRU has been powered down. If the FRU is removed without having been powered down, the 3926 mode may reboot, and the FRU may be damaged. For detailed instructions, see the chapter titled “Removable module installation” in *3926 Service Delivery Switch Installation and Startup.*

If you intend to install a Combo TDM OC48 FRU module, you can set the module type and then configure the ports before you actually insert the module in the FRU slot.

If you intend to switch from one module type to another, any provisioning specific to the old module must be deleted.

For example, to switch from TDM to NFV, all TDM FRU provisioning must be deleted. The reason for this is that auto-created TDM objects such as the TDM ports (tdm01 to tdm16) along with the port PMs are deleted when the FRU type is changed to NFV. For this reason, any provisioning against the TDM FRU ports must be deleted.

Similarly, to switch from a Combo TDM OC48 FRU module to a different type of FRU, any provisioning specific to the old module must be deleted. There are auto-created objects such as ports 7 to 16, and the user-configured OCn port or ports.

For the commands for deleting the provisioning specific to the old module, refer to *39XX/51XX Service Delivery, Aggregation, and Virtualization Switches Command Reference.*

Copyright© 2022 Ciena® Corporation

Step	Action
1	Set the FRU type for the module: `module set fru-type <NFV | TDM | OC48-TDM>` where fru-type <NFV | TDM | OC48-TDM> is the FRU type for hot swap. • NFV: The NFV module can be swapped out and replaced if this value is set. • TDM: is the default FRU type. The TDM module can be swapped out and replaced if this value is set. • OC48-TDM: The Combo TDM OC48 FRU module can be swapped out and replaced if this value is set.

Copyright© 2022 Ciena® Corporation

Procedure 10-2
Configuring TDM services

Configure TDM services to set the mode that the TDM port operates in and to enable TDM ports. By default, TDM ports are disabled.

Prior to changing the mode that the TDM port operates in, all TDM virtual circuits and associated cross connects as well as TDM profiles, attachment circuits, and associated performance monitoring instances must be deleted. For the commands, refer to *39XX/51XX Service Delivery, Aggregation, and Virtualization Switches Command Reference*.

Note 1: It takes a few seconds to re-initialize the TDM hardware to operate in the new TDM mode. During that re-initialization, TDM-related configuration changes and some TDM show commands are denied.

Note 2: On the TDM FRU, supported TDM SFPs can be used to create TDM PWs in either ANSI or ETSI mode. The combination of the mode and the TDM port type and channelization determines which SFP is supported. For example, if you are in ETSI mode and create an OC3-STM1 TDM port, then the port is an STM1 port. If the mode is ANSI, then the OC3-STM1 TDM port is an OC3 port.

Note 3: On theCombo TDM OC48 FRU module, changing the tdm mode from ANSI to ETSI or from ETSI to ANSI causes the FRU module to reboot.

Note 4: On the Combo TDM OC48 FRU module, ETSI tdm mode is supported from SAOS 6.20 onwards. If a Combo TDM OC48 FRU module is set to ETSI tdm mode, and if you downgrade to a release prior to SAOS 6.20, the downgrade changes the FRU type from OC48-TDM to TDM.

Note 5: If you downgrade to a release prior to SAOS 6.21, adaptive clock mode gets deleted from E1 and T1 SFPs port configuration.

Note 6: If you downgrade to a release prior to SAOS 6.21, all TDM configuration gets deleted for 3903, 3903x, 3906, and 3924-901(non-sync variant) platforms.

Step	Action

Copyright© 2022 Ciena® Corporation

1 Set the mode that the TDM port operates in:

```
port tdm set mode <etsi | ansi>
```

where

mode <etsi | ansi> is the mode that the TDM port operates in.

- etsi: sets the mode to European Telecommunications Standards Institute (ETSI) mode. Use ETSI mode for ports that are configured as E1.
- ansi: sets the mode to American National Standards Institute (ANSI) mode. Use ANSI mode for TSD SFP ports.

2 Enable TDM ports:

```
port tdm enable port <TDM port list>
```

—end—

Copyright© 2022 Ciena® Corporation

Example

This example shows the process of changing the tdm mode of the OC48 TDM Combo FRU Module from ANSI to ETSI.

```
> port tdm set mode etsi
WARNING: Mode changed to etsi.
Rebooting OC48-TDM FRU.
```

When the change in tdm mode begins, the FRU reboots. While the reboot is in progress, the output of the “module show” command displays the messages, “CLI command ‘module reboot’ is running,” and “module initialization is running.”

```
te3926-3231*> module show
+------------------- MODULE DETAILS-------------------+

.
.
.

+------------------------ OC48_TDM MODULE STATUS ------------------------+
| Software Status   | CLI command "module reboot" is running             |
| Hardware Status   | module initialization is running                   |
| Hardware State    | rc.combo   combo TDM FRU stopped                   |
| Hw Error Code     | 0                                                  |
| Removal Status    | NOT SAFE TO REMOVE                                 |
+------------------+-----------------------------------------------------+
```

When the reboot has finished, the FRU must re-initialize. During this time, the output of the “module show module OC48-TDM” command continues to display the message, “module initialization is running.”

```
> module show module OC48-TDM
+------------------- MODULE STATUS--------------------+
.
.
.

+------------------------ OC48_TDM MODULE STATUS ------------------------+
| Software Status   | -                      |
| Hardware Status   | module initialization is running                   |
| Hardware State    | rc.combo   combo TDM FRU started                   |
| Hw Error Code     | 0                                                  |
| Removal Status    | NOT SAFE TO REMOVE                                 |
+------------------+-----------------------------------------------------+
```

When the initialization has completed, the output of the “module show module OC48-TDM” command no longer displays the message: “module initialization is running.”

Copyright© 2022 Ciena® Corporation

Procedure 10-3
Creating TDM SFP ports

Create TDM ports and configure the related port-specific parameters for SFP TDM ports according to the network plan. When you create a TDM port, you must identify a the physical port number of the port that the SPF is inserted in.

TDM ports can be configured as

- OC3-STM1
- OC3-STM1-CEM
- OC12-STM4
- OC12-STM4-CEM
- OC48-STM16-CEM
- DS3-E3
- T1
- E1

Note 1: The “CEM” ports are created on the Combo TDM OC48 FRU module.

Note 2: On the Combo TDM OC48 FRU module, if the module is configured with ANSI mode, the DS3 and T1 ports are auto-created. If the module is configured in ETSI mode, the E3 and E1 ports are auto-created.

Copyright© 2022 Ciena® Corporation

Step	Action

To create an OC3-STM1 TDM port

1 Create an OC3-STM1 TDM port:

```
port tdm oc3-stm1 create oc3-stm1 <String[15]> port
<physical-port-number> [mac <port-mac | sfp-factory-mac>]
```

where

oc3-stm1<String[15]>	is the TDM port name. The characters :/!%,?* are not permitted in the TDM port name.
port <physical-port-number>	is the physical port number associated with the TDM port.
[mac <port-mac | sfp-factory-mac>]	is the source MAC address associated with the TDM port. Selecting the port MAC address allows an SFP to be replaced without configuration changes. The port MAC address is also required when 1+1 linear APS is supported. (See the examples in the procedure "Enabling 1+1 linear APS groups" on page 10-175.) ***Note 1:*** When using the SFP factory MAC address, a range of four is supported. When using the port MAC address, it must be an exact match. ***Note 2:*** The port-mac option is blocked if the functionality is incompatible with the transceiver type. See "Transceiver-related limitations".

To create an OC3-STM1-CEM TDM port on a Combo TDM OC48 FRU module

2 Create an OC3-STM1-CEM TDM port on a Combo TDM OC48 FRU module:

```
port tdm oc3-stm1-cem create oc3-stm1-cem <String[15]>
port <physical-port-number>
```

where

oc12-stm4-cem <String[15]>	is the TDM port name. The characters :/!%,?* are not permitted in the TDM port name.
port <physical-port-number>	is the physical port number associated with the TDM port. In this case the port number can be 2.3, 2.4, 2.5, or 2.6.

Copyright© 2022 Ciena® Corporation

To create an OC12-STM4 TDM port

3 Create an OC12-STM4 TDM port:

```
port tdm oc12-stm4 create oc12-stm4 <String[15]> port
<physical-port-number> [mac <port-mac | sfp-factory-mac>]
```

where

oc12-stm4 <String[15]>	is the TDM port name. The characters :/!%,?* are not permitted in the TDM port name.
port <physical-port-number>	is the physical port number associated with the TDM port.
[mac <port-mac | sfp-factory-mac>]	is the source MAC address associated with the TDM port. Selecting the port MAC address allows an SFP to be replaced without configuration changes. The port MAC address is also required when linear APS 1+1 APS is supported. ***Note 1:*** When using the SFP factory MAC address, a range of four is supported. When using the port MAC address, it must be an exact match. ***Note 2:*** The port-mac option is blocked if the functionality is incompatible with the transceiver type. See "Transceiver-related limitations".

To create an OC12-STM4-CEM TDM port on a Combo TDM OC48 FRU module

4 Create an OC12-STM4-CEM TDM port on a Combo TDM OC48 FRU module:

```
port tdm oc12-stm4-cem create oc12-stm4-cem <String[15]>
port <physical-port-number>
```

where

oc12-stm4-cem <String[15]>	is the TDM port name. The characters :/!%,?* are not permitted in the TDM port name.
port <physical-port-number>	is the physical port number associated with the TDM port. In this case the port number can be 2.3, 2.4, 2.5, or 2.6.

To create an OC48-STM16 TDM port

5 Create an OC48-STM16 TDM port:

```
port tdm oc48-stm16 create oc48-stm16 <String[15]> port
<physical-port-number>
```

where

oc48-stm16 <String[15]>	is the TDM port name. The characters :/!%,?* are not permitted in the TDM port name.
port <physical-port-number>	is the physical port number associated with the TDM port.

Copyright© 2022 Ciena® Corporation

To create an OC48-STM16-CEM TDM port on a Combo TDM OC48 FRU module

6 Create ***OC48-STM16-CEM TDM port***-CEM TDM port on a Combo TDM OC48 FRU module:

```
port tdm oc48-stm16-cem create oc48-stm16-cem
<String[15]> port <physical-port-number>
```

where

oc12-stm4-cem <String[15]>	is the TDM port name. The characters :/!%,?* are not permitted in the TDM port name.
port <physical-port-number>	is the physical port number associated with the TDM port. In this case the port number must be 2.3.

To create a DS3-E3 TDM port

7 Create a DS3-E3 TDM port:

```
port tdm ds3-e3 create ds3-e3 <String[15]> port
<physical-port-number> [mac <port-mac | sfp-factory-mac>]
```

where

ds3-e3 <String[15]>	is the TDM port name. The characters :/!%,?* are not permitted in the TDM port name.
port <physical-port-number>	is the physical port number associated with the TDM port.
[mac <port-mac | sfp-factory-mac>]	is the source MAC address associated with the TDM port. Selecting the port MAC address allows an SFP to be replaced without configuration changes. The port MAC address is also required when linear APS 1+1 APS is supported. ***Note:*** When using the SFP factory MAC address, a range of four is supported. When using the port MAC address, it must be an exact match.

Note: On the Combo TDM OC48 FRU module, if the module is configured in ANSI mode, DS3 ports are auto-created. Alternatively, if the module is configured in ETSI mode, E3 ports are auto-created. The names of the auto-created ports are tdm07, tdm08, tdm09, and tdm10 (corresponding to the ports labeled 13, 14, 15, and 16 on the module). The source MAC is the MAC address of the FRU.

Copyright© 2022 Ciena® Corporation

To create an E1 TDM port

8 Create an E1 TDM port:

```
port tdm e1 create e1 <String[15]> port
<physical-port-number> [mac <port-mac | sfp-factory-mac>]
```

where

e1 <String[15]>	is the TDM port name. The characters :/!%,?* are not permitted in the TDM port name.
port <physical-port-number>	is the physical port number associated with the TDM port.
[mac <port-mac | sfp-factory-mac>]	is the source MAC address associated with the TDM port. Selecting the port MAC address allows an SFP to be replaced without configuration changes. The port MAC address is also required when linear APS 1+1 APS is supported. ***Note:*** When using the SFP factory MAC address, a range of four is supported. When using the port MAC address, it must be an exact match.

Note: On the Combo TDM OC48 FRU module, if the module is configured in ETSI mode, E1 ports are auto-created. The port names are tdm01, tdm02, tdm03, tdm04, tdm05, and tdm06 (corresponding to the ports labeled 7, 8, 9, 10, 11,and 12 on the module). The source MAC is the MAC address of the FRU.

Copyright© 2022 Ciena® Corporation

To create an T1 TDM port

9 Create an T1 TDM port:

```
port tdm t1 create t1 <String[15]> port
<physical-port-number> [mac <port-mac | sfp-factory-mac>]
```

where

t1 <String[15]>	is the TDM port name. The characters :/!%,?* are not permitted in the TDM port name.
port <physical-port-number>	is the physical port number associated with the TDM port.
[mac <port-mac | sfp-factory-mac>]	is the source MAC address associated with the TDM port. Selecting the port MAC address allows an SFP to be replaced without configuration changes. The port MAC address is also required when linear APS 1+1 APS is supported. ***Note:*** When using the SFP factory MAC address, a range of four is supported. When using the port MAC address, it must be an exact match.

Note: On the Combo TDM OC48 FRU module, if the module is configured in ANSI mode, T1 ports are auto-created. The port names are tdm01, tdm02, tdm03, tdm04, tdm05, and tdm06 (corresponding to the ports labeled 7, 8, 9, 10, 11,and 12 on the module). The source MAC is the MAC address of the FRU.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-4
Deleting TDM SFP ports

Delete these TDM ports:

- OC3-STM1
- OC3-STM1-CEM
- OC12-STM4
- OC12-STM4-CEM
- OC48-STM16-CEM
- DS3-E3
- T1
- E1

Note 1: E1-TDM and T1-TDM ports are pre-created for the FRU. Only SFP TDM E1 and SFP TDM T1 ports can be deleted.

Note 2: On the Combo TDM OC48 FRU module, T1 and DS3 ports or E1 and E3 ports are auto created, depending on whether the module is configured in ANSI or ETSI mode. The auto-created ports cannot be deleted.

Step	Action

To delete an OC3-STM1 TDM port

1 Delete an OC3-STM1 TDM port:

```
port tdm oc3-stm1 delete oc3-stm1 <String[15]>
```

where

oc3-stm1<String[15]>	is the TDM port name you are deleting.

To delete an OC3-STM1-CEM TDM port on a Combo TDM OC48 FRU module

2 Delete an OC3-STM3-CEM TDM port on a Combo TDM OC48 FRU module:

```
port tdm oc3-stm1-cem delete oc3-stm1-cem <String[15]>
```

where

oc3-stm1-cem <String[15]>	is the TDM port name you are deleting.

Copyright© 2022 Ciena® Corporation

To delete an OC12-STM4 TDM port

3 Delete an OC12-STM4 TDM port:

```
port tdm oc12-stm4 delete oc12-stm4 <String[15]>
```

where

oc12-stm4 <String[15]> is the TDM port name you are deleting.

To delete an OC12-STM4-CEM TDM port on a Combo TDM OC48 FRU module

4 Delete an OC12-STM4-CEM TDM port on a Combo TDM OC48 FRU module:

```
port tdm oc12-stm4-cem delete oc12-stm4-cem <String[15]>
port <physical-port-number> [mac sfp-factory-mac]
```

where

oc12-stm4-cem <String[15]> is the TDM port name you are deleting.

To delete an OC48-STM16 TDM port

5 Delete an OC48-STM16 TDM port:

```
port tdm oc48-stm16 delete oc48-stm16 <String[15]>
```

where

oc48-stm16 <String[15]> is the TDM port name. The characters :/!%,?* are not permitted in the TDM port name.

To delete an OC48-STM16-CEM TDM port on a Combo TDM OC48 FRU module

6 Delete an OC48-STM16-CEM TDM port on a Combo TDM OC48 FRU module:

```
port tdm oc48-stm16-cem delete oc48-stm16-cem
<String[15]>
```

where

oc48-stm16-cem <String[15]> is the TDM port name you are deleting.

To delete a DS3-E3 TDM port

7 Delete a DS3-E3 TDM port:

```
port tdm ds3-e3 create ds3-e3 <String[15]>
```

where

ds3-e3 <String[15]> is the TDM port name you are deleting.

To delete an E1 TDM port

8 Delete an E1 TDM port:

```
port tdm e1 create e1 <String[15]>
```

where

e1 <String[15]> is the TDM port name you are deleting.

Copyright© 2022 Ciena® Corporation

To delete an T1 TDM port

9 Delete a T1 TDM port:

```
port tdm t1 delete t1 <String[15]>
```

where

t1 <String[15]> is the TDM port name you are deleting.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-5
Loading a new firmware load on an SFP

When a new shipment of SFPs is received, they may not have the correct firmware load. In this case, each SFP requires an upgrade immediately. This procedure can be done at any time, not just after an upgrade.

Note: An SFP must be created and enabled to perform this procedure.

Step	Action
1	Determine if the upgrade is compatible by looking at the “Upgrade Operation” field: `software xcvr port`
2	Start the upgrade of the SFP firmware: `software upgrade port <portList> [now]` where port <portList> — is the list of one or more SFP port numbers you are loading the new ISD load on.
3	Check the progress of the SFP firmware upgrade by looking at the “Upgrade Operation” field: `software show port <portList>` where port <portList> — is the list of one or more SFP port numbers.

—end—

Example

This example shows the output of the “software show port” command.

```
> software show port 1
+--------------- XCVR DATA - Port 1 ---------------+
| Parameter                   | Value               |
+-----------------------------+---------------------+
| Running Firmware Version    | 03003200            |
| Available Firmware Version  | 03003200            |
| Firmware Operation Require  | None                |
| Firmware Update Result      | Successful          |
+-----------------------------+---------------------+
| Running Bank                | A                   |
| Running Version             | 0802000001          |
| Running Date                | 03 Apr 17 13:29     |
+-----------------------------+---------------------+
| Available Version           | 0802000001          |
+-----------------------------+---------------------+
| Operation Required          | None                |
| Version Compatibility       | Yes                 |
```


Copyright© 2022 Ciena® Corporation

```
| Traffic Impacted            | No                  |
| Update Result               | Successful          |
| Update Load Progress        | 0  %                |
+-----------------------------+---------------------+
```

The command indicates whether the upgrade is compatible or not. This table describes the parameters from the “upgrade xcvr show command”.

Table 10-23
Upgrade xcvr show command

Parameter	Explanation
Running Firmware Version	Hardware reversion of the SFP. When an SFP needs to be downgraded, this field indicates “Not available”.
Available Firmware Version	Firmware version available in SAOS for upgrade/downgrade.
Firmware Operation Require	An upgrade is available. If the SFP is up to date, no upgrade is available.
Running Bank	Current active firmware bank. Values are A or B. The SFP upgrade writes to the non-active bank which becomes active after the upgrade is completed. When an SFP needs to be downgraded, this field indicates “Not available”.
Running Version	Firmware version on the running bank. When an SFP needs to be downgraded, this field indicates “Not available”.
Running Date	Date the firmware was built on the running back. When an SFP needs to be downgraded, this field indicates “Not available”.
Available Version	Firmware version available in SAOS for upgrade/downgrade.
Operation Required	An upgrade is available. If the SFP is up to date, no upgrade is available.

Copyright© 2022 Ciena® Corporation

Table 10-23
Upgrade xcvr show command

Parameter	Explanation
Version Compatibility	Yes means this is a minor upgrade, and there is minimum traffic impact when the SFP upgrades. This typically means less than three seconds. No means it is a major upgrade and traffic does not flow until the SFP is upgraded. When no is indicated, it means that the SFP may also be deleted/re-created (internally) as part of the upgrade. This means that the SFP may temporarily appear as not inserted. It reappears in 10 seconds.
Traffic Impacted	Indicates if the traffic impacted by the upgrade.
Update Result	Status of the upgrade.
Update Load Progress	Indicates the percentage of firmware loaded in increments of 1 percent once the upgrade has started.

This example shows a command to upgrade multiple ports.

```
> so upgrade port 1-24 now
PRESS <CTRL>C TO ABORT   (10 seconds)
           PRESS <CTRL>C TO ABORT   ( 9 seconds)
           PRESS <CTRL>C TO ABORT   ( 8 seconds)
           PRESS <CTRL>C TO ABORT   ( 7 seconds)
           PRESS <CTRL>C TO ABORT   ( 6 seconds)
           PRESS <CTRL>C TO ABORT   ( 5 seconds)
           PRESS <CTRL>C TO ABORT   ( 4 seconds)
           PRESS <CTRL>C TO ABORT   ( 3 seconds)
           PRESS <CTRL>C TO ABORT   ( 2 seconds)
           PRESS <CTRL>C TO ABORT   ( 1 seconds)
NOTE: Multiple TDM SFP specified, preforming bulk upgrade.
      Operation takes around 5 min per SFP and can only be interrrupted if stop command is issued.
```

This example shows a bulk “software show port” command.

```
so sh port 1-24
+-------------------------------------------------------------------------------------------------+
|              TDM SFP Ports Bulk Upgrade Status                                                  |
+----------+------+--------- +-----------+-----------+-----------+-----------+----------+---------+
| TDM      | Phys | Admin    | Available | Running   | Operation | Upgrade   | Load     | Bulk    |
| Port     | Port | State    | Version   | Version   | Required  | Result    | Progress | Status  |
+----------+------+--------- +-----------+-----------+-----------+-----------+----------+---------+
|  tsop1_9 |    9 | disabled | Not Avail | Not Avail |      None | Successful |     0% | Not Run |
| csop1_10 |   10 | disabled | Not Avail | Not Avail |      None | Successful |     0% | Not Run |
| csop1_11 |   11 | disabled | Not Avail | Not Avail |      None | Successful |     0% | Not Run |
| tsop2_14 |   14 | disabled | Not Avail | Not Avail |      None | Successful |     0% | Not Run |
| cep2_15  |   15 | disabled | Not Avail | Not Avail |      None | Successful |     0% | Not Run |
| cep2_16  |   16 | disabled | Not Avail | Not Avail |      None | Successful |     0% | Not Run |
| tsop3_22 |   22 | disabled | Not Avail | Not Avail |      None | Successful |     0% | Not Run |
| tsop3_23 |   23 | disabled | Not Avail | Not Avail |      None | Successful |     0% | Not Run |
+----------+------+--------- +-----------+-----------+-----------+-----------+----------+---------+
```


Copyright© 2022 Ciena® Corporation

Copyright© 2022 Ciena® Corporation

Procedure 10-6
Creating and enabling a port to be used by a TSOP transceiver

Use this procedure to create and enable a port to be used by a TSOP transceiver.

Step	Action

To create an OC3 or STM1 port to be used by a TSOP01 transceiver

1 Create the OC3 or STM1 port.

```
port tdm oc3-stm1 create oc3-stm1 <port_name> port <port-number>
```

where

<port-name>	is the name assigned to the port.
<port-number>	is the physical port number.

Enable the OC3 or STM1 port

2 Enable the OC3 or STM1 port.

```
port tdm enable port <port-name>
```

To create an OC12 or STM4 port to be used by a TSOP01 transceiver

3 Create the OC12 or STM4 port.

```
port tdm oc12-stm4 create oc12-stm4 <port_name> port <port-number>
```

where

<port-name>	is the name assigned to the port.
<port-number>	is the physical port number.

Enable the OC12 or STM4 port

4 Enable the OC12 or STM4 port.

```
port tdm enable port <port-name>
```

To create an OC48 or STM16 port to be used by a TSOP01 transceiver

5 Create the OC48 or STM16 port.

```
port tdm oc48-stm16 create oc48-stm16 <port_name> port <port-number>
```

where

<port-name>	is the name assigned to the port.
<port-number>	is the physical port number.

Copyright© 2022 Ciena® Corporation

Enable the OC48 or STM16 port

6 Enable the OC48 or STM16 port.

```
port tdm enable port <port-name>
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-7
Changing the section trace

Change the section trace for a TDM port when device ID information changes.

The section-trace attribute sets the SONET Overhead J0 byte used to communicate device ID information. The SONET Overhead J0 byte is transmitted by the equipment where the SDH or SONET signal originates and persists through the network to the equipment where the SDH or SONET signal terminates. The device ID information is used by the receiving equipment to verify its continued connection to the intended transmitter.

Step	Action
1	Change the section trace: `port tdm <oc3-stm1|oc12-stm4|ds3-e3> set <oc3-stm1|oc12-stm4|ds3-e3> <tdm-port-name> [section-trace <String[15]>]` where <oc3-stm1 | oc12-stm4 | ds3-e3> <tdm-port-name> — is the TDM port name. section-trace<String[15]>] — is a user-defined 15-character value. The default value of this attribute is the sequence “01” repeated to create a string 15 characters long, which is interpreted as no section-trace value.

To change the section trace on a port on a Combo TDM OC48 FRU module

2	Change the section trace on a port on a Combo TDM OC48 FRU module: `port tdm <oc3-stm1-cem|oc12-stm4-cem|oc48-stm16-cem> set <oc3-stm1-cem|oc12-stm4-cem|oc48-stm16-cem> <tdm-port-name> [section-trace <String[15]>]` where <oc3-stm1-cem | oc12-stm4-cem | oc48-stm16-cem> <tdm-port-name> — is the TDM port name. section-trace<String[15]>] — is a user-defined 15-character value. The default value of this attribute is the sequence “01” repeated to create a string 15 characters long, which is interpreted as no section-trace value.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-8
Configuring a protection group

For the 3926 TDM ports, a protection group specifies the selection of input timing reference.

Step	Action
1	Configure a protection group on a 3926: `port tdm <t1/e1> set port <tdm port name> protection-group <group name>` where port <tdm port name> — is the TDM port name. <t1/e1> — T1 is only available when the TDM mode is ANSI. E1 is only available when the TDM mode is ETSI. ***Note:*** You need to change the TDM mode to ETSI to see E1 commands. protection-group <group-name> — is the protection-group name. **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 10-9
Configuring a line-code on a T1 module

For the T1 3926 TDM ports line encoding must be specified.

Step	Action
1	Configure a line code on a T1 on a 3926: `port tdm t1 set port <tdm port name> line-code <b8zs|ami>` where port tdm <tdm port name> — is the TDM port name. line-code <b8zs|ami> — specifies bipolar 8 zero substitution (b8zs) or alternative mark inversion (ami). Default is b8zs. **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 10-10
Configuring a line-code on an E1 module

For the E1 3926 TDM ports line encoding must be specified.

Step	Action
1	Configure a line code on an E1 on a 3926: `port tdm e1 set port <tdm port name> line-code <hdb3|ami>` where port tdm <tdm port name> — is the TDM port name. line-code <hdb3|ami> — specifies high density bipolar order 3 (hdb3) or alternative mark inversion (ami). Default is hdb3. **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 10-11
Configuring framing on a T1 module

For the T1 3926, TDM ports framing must be specified.

Step	Action
1	Configure framing on a T1 on a 3926: `port tdm t1 set port <tdm port name> framing <unframed|super-frame|extended-super-frame>` where port tdm <tdm port name> — is the TDM port name. framing <unframed|super-frame|extended-super-frame> — specifies the framing type. Default is unframed. **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 10-12
Configuring framing on a E1 module

For the E1 3926, TDM ports framing must be specified.

Step	Action
1	Configure framing on a E1 on the 3926: `port tdm e1 set port <tdm port name> framing <unframed|basic|crc-multiframe>` where port tdm <tdm port name> — is the TDM port name. framing <unframed|basic|crc-multiframe> — specifies the framing type. Default is unframed. ***Note:*** If you have configured MEF8 virtual circuits on an E1 TDM port with SAToP services on the attachment circuit, and if you want to change the E1 framing (from unframed to framed or from framed to unframed), you must delete the MEF8 virtual circuits before changing the framing. **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 10-13
Configuring framing on a Combo TDM OC48 FRU module

Use this procedure to configure framing for the DS3 and T1 ports on a Combo TDM OC48 FRU module.

Step	Action

To configure framing for a DS3 port on a Combo TDM OC48 FRU module

1 Configure framing for the a DS3 port on a Combo TDM OC48 FRU module:

```
port tdm ds3-e3 set ds3-e3 <tdm port name>
framing <unframed | c-bit | m13>
```

where

port tdm <tdm port name>	is the port name.
framing <unframed | m13 |c-bit>	specifies the framing type. Default is unframed. ***Note:*** When DS3 is unchannelized, the framing type can be unframed or c-bit. When DS3 is channelized to DS1, the framing type can be c-bit or m13. When DS3 is channelized to STS1, only c-bit framing can be used.

To configure framing for an E3 port on a Combo TDM OC48 FRU module

2 Configure framing for an E3 port on a Combo TDM OC48 FRU module:

```
port tdm ds3-e3 set ds3-e3 <tdm port name>
framing <unframed | g751>
```

where

port tdm <tdm port name>	is the port name.
framing <unframed | g751>	specifies the framing type. Default is unframed. ***Note:*** For E3 ports on the Combo TDM OC48 FRU module, loopback is supported only with g751 framing.

Copyright© 2022 Ciena® Corporation

To configure framing for a T1 port on a Combo TDM OC48 FRU module

3 Configure framing for a T1 port on a Combo TDM OC48 FRU module:

```
port tdm t1 set t1 port <tdm port name>
framing <unframed | esf | sf>
```

where

port tdm <tdm port name>	is the port name.
framing <unframed | esf | sf>	specifies the framing type. Default is unframed.

To configure framing for an E1 port on a Combo TDM OC48 FRU module

4 Configure framing for an E1 port on a Combo TDM OC48 FRU module:

```
port tdm e1 set port <tdm port name>
framing <unframed | basic | crc-multiframe>
```

where

port tdm <tdm port name>	is the port name.
framing <unframed | basic | crc-multiframe>	specifies the framing type. Default is unframed. ***Note 1:*** For E3 ports on the Combo TDM OC48 FRU module, loopback is supported only with g751 framing. ***Note 2:*** If you have configured MEF8 virtual circuits on an E1 TDM port with SAToP services on the attachment circuit, and if you want to change the E1 framing (from unframed to framed or from framed to unframed), you must delete the MEF8 virtual circuits before changing the framing.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-14
Configuring SSM

SSM on the E1 FRU

Synchronization status messages (SSM) can be enabled or disabled for E1 field replaceable unit (FRU) ports. When enabled, SSM contributes to the clock selection algorithm. SSM is disabled by default.

The spare bit to use for SSM, from sa4 to sa8, can also be configured.

SSM on the Combo TDM OC48 FRU module

On the Combo TDM OC48 FRU module, synchronization status messages (SSM) are supported, but SSM is not user-configurable. SSM behavior is based on the clock mode, as follows:

- Internal: SSM Tx value is provide by the primary or secondary port reference (managed by Sync)
- Line: SSM Tx value is send based on SSM Rx received on the port.
- ACR/DCR: SSM Tx is currently hard-coded to PRC.

Note 1: The Combo TDM OC48 FRU module supports SSM Rx on DS1 ports, with all clock modes.

Note 2: The Combo TDM OC48 FRU module SSM Rx on OCn ports.

Note 3: The Combo TDM OC48 FRU module supports SSM Tx on OCn ports only with internal clock mode.

Step	Action

To configure SSM on an E1 module port

1 Configure SSM on an E1 module port:

```
port tdm e1 ssm <enable | disable> e1 <tdm port name>
```

where

port tdm <tdm port name>	is the TDM port name.
ssm <enable | disable>	specifies whether SSM is enabled. Default is disable.

Copyright© 2022 Ciena® Corporation

To configure the spare bit to use for SSM

2 Configure the spare bit to use for SSM:

```
port tdm e1 ssm set e1 <tdmPortName> ssm-location <sa4 |
sa5 | sa6 | sa7 | sa8>
```

where

port tdm <tdm port name>	is the TDM port name.
ssm-location <sa4 | sa5 | sa6 | sa7 | sa8>	specifies the spare bit location to use for SSM. Default is sa8.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-15
Configuring loopcode transmission on T1 module ports

There are two kinds of loopcode control signaling, as summarized in this comparison table:

Table 10-24
Loopcode control signaling types

Attributes	FDL (facility data link)	In-band
Frame type	Should be extended super frame (ESF).	Should be super frame (SF).
Supported loopcodes	• Line activate • Line deactivate • Network loop activate • Universal deactivate • Loopback retention • Idle	• Line activate • Line deactivate • Network interface unit (NIU) activate • NIU deactivate

Loopcode transmission can be enabled or disabled for T1 field replaceable unit (FRU) ports.

Note 1: Loopback cannot be enabled at the same time loopcode transmission is active.

Note 2: Loopcode transmission is not supported for the Combo TDM OC48 FRU module.

Codewords accepted and allowed for loopcode transmission are:

- network-side-loopback-activate
- line-loopback-activate
- niu-loopback-activate*
- loopback-retention
- universal-loopback-deactivate
- line-loopback-deactivate
- niu-loopback-deactivate*
- idle

* is applicable only to In-band loopcode control signaling types.

Copyright© 2022 Ciena® Corporation

Once set, loopcode transmissions behave as follows:

Table 10-25
Loopcode transmission behavior

FDL (facility data link)	In-band
• The activate or deactivate loopcode is transmitted for up to 3 seconds, then the idle loopcode is sent. • The retention loopcode is transmitted until another loopcode is set or retention fails. • The idle loopcode is transmitted until another loopcode is set.	• The activate or deactivate loopcode is transmitted for 12 seconds. • After 5 seconds, the verification period begins and the loopcode must be confirmed before the 12 seconds have expired.

Loopcode transmission set is not persistent (not saved to the configuration). Also, loopcode transmission cannot be set if the port is disabled.

Step	Action

To configure loopcode transmission on a T1 module port

1 Configure loopcode transmission on a T1 module port:

```
port tdm t1 loopcode <enable | disable> t1 <tdm port name>
```

where

port tdm <tdm port name>	is the TDM port name.
loopcode <enable | disable>	specifies whether loopcode transmission is enabled. Default is disable.

To configure the codeword for loopcode transmission on a T1 module port

2 Configure the codeword for loopcode transmission on a T1 module port:

```
> port tdm t1 loopcode set t1 <tdmPortName> codeword <network-side-loopback-activate | line-loopback-activate | niu-loopback-activate | loopback-retention | universal-loopback-deactivate | line-loopback-deactivate | niu-loopback-deactivate | idle>
```

where

port tdm <tdm port name>	is the TDM port name.
codeword <network-side-loopback-activate | line-loopback-activate | niu-loopback-activate | loopback-retention | universal-loopback-deactivate | line-loopback-deactivate | niu-loopback-deactivate | idle>	specifies the loopcode codeword to use.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-16
Configuring C-TPID on the 3926 T1/E1 port

Configure C-TPID on the 3926 T1/E1 port and return the C-TPID to default values (0x8100).

> ***Note:*** For the Combo TDM OC48 FRU module, the default C-TPID is 0x8100 and does not use the global C-TPID configuration. The only way to change the C-TPID is via the C-TPID attribute in the virtual-circuit tdm vc object.

Step	Action

To configure C-TPID on the 3926 T1/E1 port

1 Configure C-TPID on the 3926 T1/E1 port:

```
port tdm set c-tpid <8100|9100|88a8>
```

where

c-tpid <8100|9100|88a8>	sets the customer TPID on the 3926.

To return C-TPID on the 3926 T1/E1 port to default values

2 Return C-TPID on the 3926 T1/E1 port to default values:

```
port tdm unset c-tpid
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-17
Configuring C-TPID on the TDM SFP

Configure C-TPID on the TDM SFP.

Use this option to override the global default C-TPID.

> ***Note:*** For the Combo TDM OC48 FRU module, the default C-TPID is 0x8100 and does not use the global C-TPID configuration. The only way to change the C-TPID is by the c-tpid attribute of the virtual circuit, as covered in the procedure "Configuring TDM virtual circuits".

Step	Action

To configure C-TPID for TDM SFP

1 Configure C_TPID on the TDM SFP:

```
port tdm oc3-stm1/oc12-stm4/ds1-e1/ds3-e3 set port <tdm
port name> c-tpid <8100|9100|88a8>
```

where

oc3-stm1/oc12-stm4/ds1-e1/ds3-e3	is the oc3-stm1/oc12-stm4/ds1-e1/ds3-e3 port type.
<tdm port name>	is the TDM port name.
c-tpid <8100|9100|88a8>	sets the C-TPID on the TDM SFP.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-18
Configuring the cable length

For DS3 TDM ports, a cable length must be specified. The choices are normal and extended where normal is 0 to 225 feet and extended is 225 feet to 400 feet. The default is normal.

Note: Cable length is not supported for E3.

Step	Action
1	Configure the cable length: `port tdm ds3-e3 set ds3-e3 <ds3> [cable-length <normal|extended>]` where ds3-e3 <ds3> — is the ds3 port name. cable-length <normal | extended> — sets the cable length. Normal is 0..225FT. Extended is 225..450FT. Default is normal **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 10-19
Configuring the line build out cable length

For T1 ports on the TDM FRU module, and for T1 ports on the Combo TDM OC48 FRU module, the line build out cable length must be specified. The choices are 133, 266, 399, 533 and 655 feet. For TPoP DS1 ports, the cable length (line build out) must also be specified. The default is 133.

Step	Action
1	Configure the line build out cable length: `port tdm t1 set port <port> line-build-out <line-build-out>` where port <port> — is the name of the port. line-build-out <line-build-out> — is the line-build-out value in feet. Values are 133, 266, 399, 533 or 655 feet. **—end—**

Example

This example configures the line-build-out cable length:

```
port tdm1 set port tdmPort1 line-build-out 655
```


Copyright© 2022 Ciena® Corporation

Procedure 10-20
Configuring TDM SFP PRBS insertion and monitoring

Configure Pseudo-Random Bit Sequence (PRBS) insertion and monitoring on TDM SFP ports.

PRBS insertion and monitoring is used as a traffic test to ensure a datapath is working properly. It is typically used during initial deployment.

PRBS insertion can be generated in the ingress direction and monitored in the egress direction. PRBS insertion and monitoring can be applied on both the line and network sides, with the exception of DS3 TDM SFP ports. Only network insertion and monitoring are available for DS3 TDM SFP ports. If PRBS is detected while the monitoring option is on, an alarm is generated.

This figure illustrates line PRBS.

Figure 10-4
Line PRBS

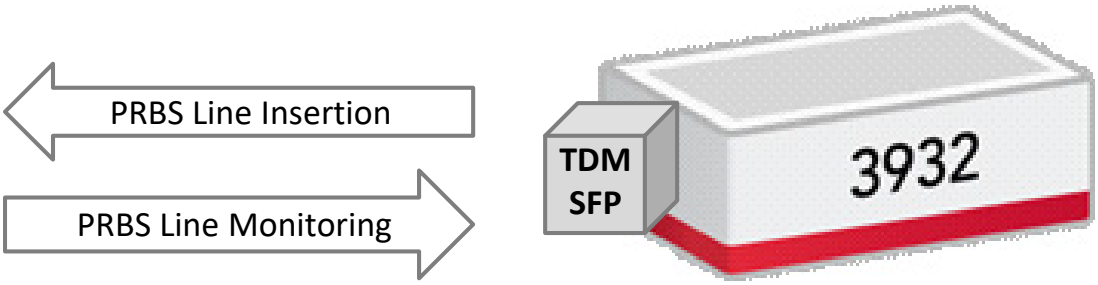

This figure illustrates network PRBS.

Figure 10-5
Network PRBS

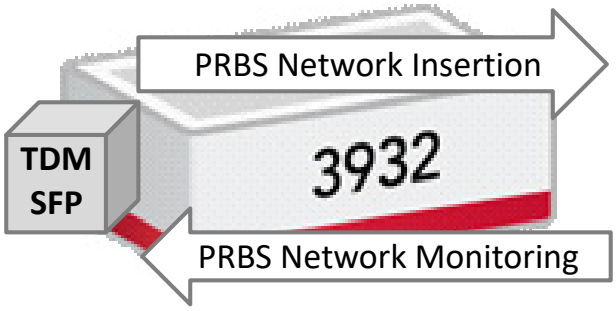

Note 1: If the PRBS direction is not the same for insertion and monitoring, the channel assigned to insertion and to monitoring must be the same.

Note 2: PRBS is not supported on the GES101 STM1 and TDM2GE E1.

Copyright© 2022 Ciena® Corporation

TDM SFP PRBS and the Combo TDM OC48 FRU module

The following limitations apply to TDM SFP PRBS on the Combo TDM OC48 FRU module:

- PRBS insertion and monitoring is not supported on the network side.
- The FRU does not support PRBS for DS3->STS1 in either direction.

Step Action

To set OC3-STM1 TDM SFP port PRBS insertion and monitoring attributes:

1 Set OC3-STM1 TDM SFP port PRBS insertion and monitoring attributes:

```
port tdm oc3-stm1 prbs set oc3-stm1 <tdm-port-name>
insert <line|network|none> ins-channel <ins-channel>
monitoring <line|network|none> mon-channel <mon-channel>
```

where

oc3-stm1 <tdm-port-name>	is the TDM port name.
insert <line | network | none>	is the side on which insertion is applied.
ins-channel <ins-channel>	is the channel on which insertion is applied.
imonitoring <line | network | none>	is the side on which monitoring is applied.
mon-channel <mon-channel>	is the channel on which monitoring is applied.

Note: This command is blocked if the functionality is incompatible with the transceiver type. See "Transceiver-related limitations".

Copyright© 2022 Ciena® Corporation

For Combo TDM OC48 FRU, to set OC3-STM1 TDM SFP port PRBS insertion and monitoring attributes:

2 For an OC3 port on the Combo TDM OC48 FRU module, set the OC3-STM1 TDM SFP PRBS insertion and monitoring attributes:

```
port tdm oc3-stm1-cem prbs set oc3-stm1-cem <tdm-port-name> insert <line|none> ins-channel <ins-channel> monitoring <line|none> mon-channel <mon-channel>
```

where

oc3-stm1-cem <tdm-port-name>	is the TDM port name.
insert <line | none>	is the side on which insertion is applied.
ins-channel <ins-channel>	is the channel on which insertion is applied.
imonitoring <line | none>	is the side on which monitoring is applied.
mon-channel <mon-channel>	is the channel on which monitoring is applied.

Note: This command is blocked if the functionality is incompatible with the transceiver type. See "Transceiver-related limitations".

To set OC12-STM4 TDM SFP port PRBS insertion and monitoring attributes:

3 Set OC12-STM4 TDM SFP port PRBS insertion and monitoring attributes:

```
port tdm oc12-stm4 prbs set oc12-stm4 <tdm-port-name> insert <line|network|none> ins-channel <ins-channel> monitoring <line|network|none> mon-channel <mon-channel>
```

where

oc12-stm4 <tdm-port-name>	is the TDM port name.
insert <line | network | none>	is the side on which insertion is applied.
ins-channel <ins-channel>	is the channel on which insertion is applied.
imonitoring <line | network | none>	is the side on which monitoring is applied.
mon-channel <mon-channel>	is the channel on which monitoring is applied.

Note: This command is blocked if the functionality is incompatible with the transceiver type. See "Transceiver-related limitations".

Copyright© 2022 Ciena® Corporation

For Combo TDM OC48 FRU, to set OC12-STM4 TDM SFP port PRBS insertion and monitoring attributes:

4 For an OC12 port on the Combo TDM OC48 FRU module, set the OC12-STM4 TDM SFP PRBS insertion and monitoring attributes:

```
port tdm oc12-stm4-cem prbs set oc12-stm4-cem <tdm-port-name> insert <line|none> ins-channel <ins-channel> monitoring <line|none> mon-channel <mon-channel>
```

where

oc12-stm4-cem <tdm-port-name>	is the TDM port name.
insert <line | none>	is the side on which insertion is applied.
ins-channel <ins-channel>	is the channel on which insertion is applied.
imonitoring <line | network | none>	is the side on which monitoring is applied.
mon-channel <mon-channel>	is the channel on which monitoring is applied.

Note: This command is blocked if the functionality is incompatible with the transceiver type. See "Transceiver-related limitations".

For Combo TDM OC48 FRU, to set OC48-STM16 TDM SFP port PRBS insertion and monitoring attributes:

5 For an OC48 port on the Combo TDM OC48 FRU module, set the OC48-STM16 TDM SFP PRBS insertion and monitoring attributes:

```
port tdm oc48-stm16-cem prbs set oc48-stm16-cem <tdm-port-name> insert <line|none> ins-channel <ins-channel> monitoring <line|none> mon-channel <mon-channel>
```

where

oc48-stm16-cem <tdm-port-name>	is the TDM port name.
insert <line | none>	is the side on which insertion is applied.
ins-channel <ins-channel>	is the channel on which insertion is applied.
imonitoring <line | network | none>	is the side on which monitoring is applied.
mon-channel <mon-channel>	is the channel on which monitoring is applied.

Note: This command is blocked if the functionality is incompatible with the transceiver type. See "Transceiver-related limitations".

Copyright© 2022 Ciena® Corporation

To set DS3-E3 TDM SFP port PRBS insertion and monitoring attributes:

6 Set DS3-E3 TDM SFP port PRBS insertion and monitoring attributes:

```
port tdm ds3-e3 prbs set ds3-e3 <tdm-port-name> insert
<network|none> monitoring <network|none>
```

where

sd3-e3 <tdm-port-name>	is the TDM port name.
insert <network | none>	is the side on which insertion is applied.
imonitoring < network | none>	is the side on which monitoring is applied.

Note 1: PBRS monitoring is not supported for DS3-E3 ports supporting SAToP services if framing is configured.

Note 2: PRBS is not supported for DS3 ports supporting SAPoP services if the attachment-circuit channelization is DS1.

Note 3: PRBS is supported for a DS3 or E3 port supporting SAToP services if the attachment-circuit channelization type is identical to the port type (DS3 channelization for a DS3 port or E3 channelization for an E3 port), and if the framing for the port is set to “unframed”.

To set T1 TDM port PRBS insertion and monitoring attributes on T1 ports on 3926 and on the Combo TDM OC48 FRU

7 Set T1 or TDM port PRBS insertion and monitoring attributes on T1 ports on 3926 and on the Combo TDM OC48 FRU module:

```
port tdm t1 prbs set t1 <tdmPortName> insert
<line|network|none> ins-channel <channel #> monitoring
<line|network|none> mon-channel <channel #>
```

where

t1 <TDMPortName>	is the TDM port name.
insert <line|network|none>	sets PRBS insertion on a channel.
monitoring <line|network|none>	sets PRBS monitoring on a channel.
mon-channel <channel #>	is the PRBS monitoring channel number.

Note: PRBS line monitoring is not supported for a T1 ports that is supporting SAToP services if the framing for the TI port is set to “super-frame” or to “extended-super-frame”.

Copyright© 2022 Ciena® Corporation

To set E1 TDM port PRBS insertion and monitoring attributes on E1 ports on 3926 and on the Combo TDM OC48 FRU

8 Set E1 TDM port PRBS insertion and monitoring attributes on 3926 E1 ports on 3926 and on the Combo TDM OC48 FRU module:

```
port tdm e1 prbs set e1 <tdmPortName> insert
<line|network|none> ins-channel <channel #> monitoring
<line|network|none> mon-channel <channel #>
```

where

t1 <TDMPortName>	is the TDM port name.
insert <line|network|none>	sets PRBS insertion on a channel.
monitoring <line|network|none>	sets PRBS monitoring on a channel.
mon-channel <channel #>	is the PRBS monitoring channel number.

To return T1 TDM port PRBS insertion and monitoring attributes to default values:

9 Return T1 TDM port PRBS insertion and monitoring attributes to default values:

```
port tdm t1 prbs unset t1 <tdmPortName> insert monitoring
```

where

t1 <TDM name>	is the TDM port name.

To return E1 TDM port PRBS insertion and monitoring attributes to default values:

10 Return E1 TDM port PRBS insertion and monitoring attributes to default values:

```
port tdm e1 prbs unset e1 <tdmPortName> insert monitoring
```

where

e1 <TDM name>	is the TDM port name.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-21
Recovering the clock

The TDM SFPs recover the clocking in many ways. The clock recovery mode can be internal (retiming), ACR, DCR or EPAR. For OC3 and OC12, DS3 STM1 and STM4 TDM ports, the default is DCR-EPAR. For E1 ports, the default is DCR-EPAR while in non-sync platforms(3903, 3903x, 3906, 3924) the default is ACR. For SFP E1 ports, adaptive and recovered are not applicable. For PWE E1 ports, the DCR-EPAR option is not applicable.

> ***Note 1:*** Internal timing-mode is not available to OC3, OC12 and DS3 ports. Only DCR-EPAR is supported. For OC12, DS3, STM4, there is no clock-mode option as only DCR-EPAR is supported.
>
> ***Note 2:*** The TDM 3926 ports use a protection group to recover the clock. See "Configuring TDM synchronization frequency protection groups" on page 10-92.

When the TDM mode changes from ETSI to ANSI and there are OC3-STM1 ports created with clock-mode set to internal, the clock-mode must be changed to DCR-EPAR before the TDM mode change is allowed.

> ***Note:*** Internal timing mode is only applicable to STM1 ports with channelization of E1. If clock mode is set to internal, attachment circuit with channelization that is not DS1-E1 cannot be created while in ETSI mode.

ACR clock-mode supports the following three topologies:

- **Topology 1:** TDM Clock-1 acts as the master and provides clocks to all other nodes.

Figure 10-6
ACR Topology 1

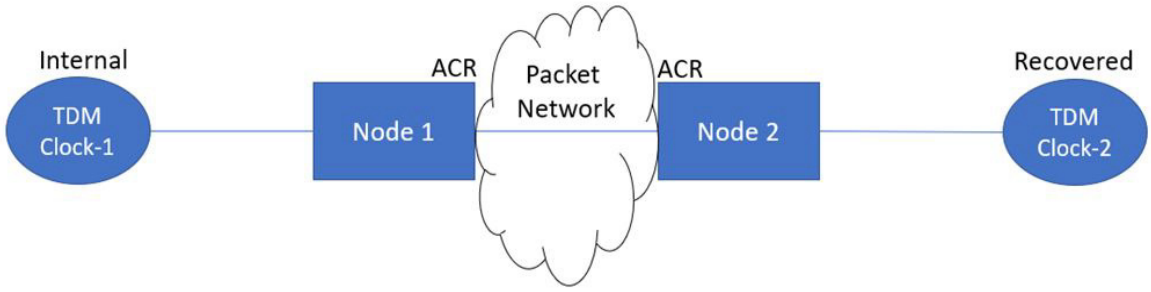

- **Topology 2:** Node-1 acts as the master and provides clocks to all other nodes.

Copyright© 2022 Ciena® Corporation

Figure 10-7
ACR Topology 2

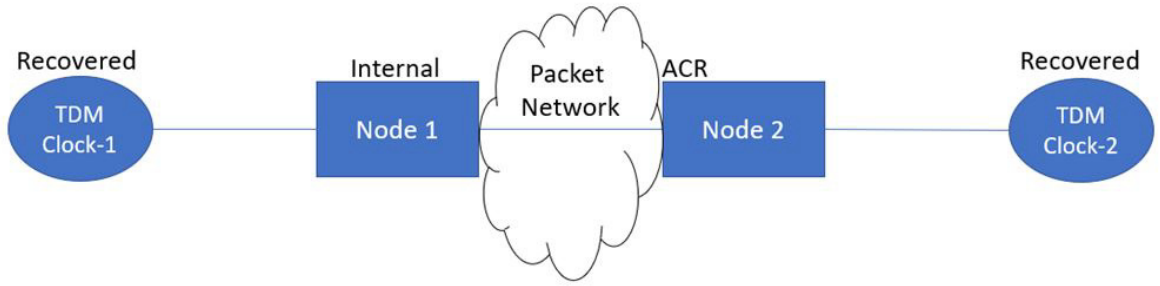

- **Topology 3:** External clock source provides a clock to all other nodes. Hence, all nodes get synchronized. If the external clock source is unavailable, then node-1 and node-2 handles different clock from TDM clock-1 and clock-2 and does not result in bit-errors. If node-1 and node-2 are unable to compensate for the clock difference of each other, then bit-errors may show on TDM Equipment emerging TDM clock-1 and clock-2.

ACR Topology 3

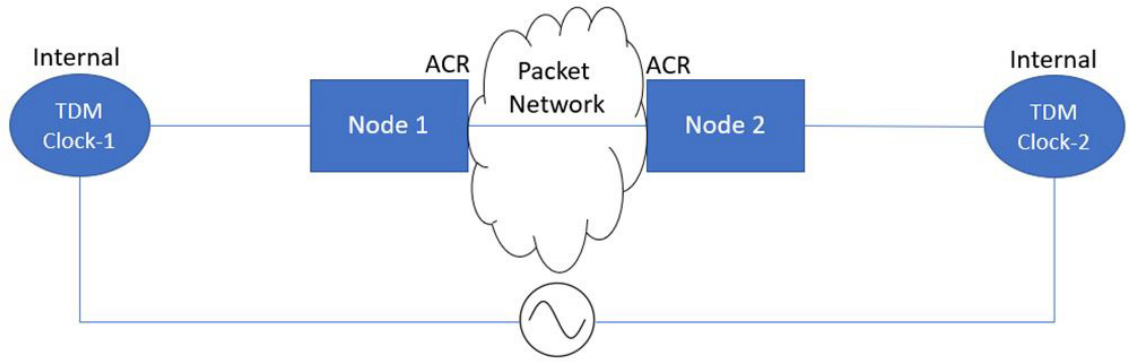

This table describes the clock modes.

Table 10-26
Clock modes

TDM port type	TDM mode	Channelization	Valid clock mode	Default clock
OC3-STM1	ANSI	STS3c, STS1	EPAR	EPAR
OC12-STM4	ANSI	STS12c, STS3c, STS1	EPAR	EPAR
OC3-STM1	ETSI	E1	DCR, INTERNAL	DCR
OC3-STM1	ETSI	VC4c	EPAR	EPAR

Copyright© 2022 Ciena® Corporation

Table 10-26
Clock modes

TDM port type	TDM mode	Channelization	Valid clock mode	Default clock
OC12-STM4	ETSI	VC4c, VC4-4c	EPAR	EPAR
DS3-E3	ANSI	STSI	EPAR	EPAR
E1 (Default Clock)	ETSI	E1	ACR (3903, 3903x, 3906, 3924-901), DCR, INTERNAL	DCR
T1 (Default Clock)	ANSI	DS1	ACR (3903, 3903x, 3906, 3924-901), DCR, INTERNAL	DCR
E1 [TDM FRU module)	ETSI	E1, E0	DCR, INTERNAL, ACR, RECOVERED	INTERNAL
T1 [TDM FRU module)	ANSI	DS1, DS0	DCR, INTERNAL, ACR, RECOVERED	INTERNAL
OC3-STM1-CEM (Combo TDM OC48 FRU)	ANSI	STS3c, STS1	EPAR, INTERNAL	INTERNAL
		DS1	DCR, INTERNAL	INTERNAL
	ETSI	VC3, VC4	EPAR, INTERNA	INTERNAL
		E1	DCR, INTERNAL	INTERNAL
OC12-STM4-CEM (Combo TDM OC48 FRU)	ANSI	STS12c, STS3c, STS1	EPAR, INTERNAL	INTERNAL
		DS1	DCR, INTERNAL	INTERNAL
	ETSI	VC4-4c, VC4, VC3	EPAR, INTERNAL	INTERNAL
		E1	DCR, INTERNAL	INTERNAL
OC48-STM16-CEM (Combo TDM OC48 FRU)	ANSI	STS48c, STS12c, STS3c, STS1	EPAR, INTERNAL	INTERNAL
	ETSI	VC4-16C, VC4-4C, VC4, VC3	EPAR, INTERNAL	INTERNAL
OC48-STM16-CEM (Combo TDM OC48 FRU)	ANSI	DS1	DCR, INTERNAL	INTERNAL
	ETSI	E1	DCR, INTERNAL	INTERNAL

Copyright© 2022 Ciena® Corporation

Table 10-26
Clock modes

TDM port type	TDM mode	Channelization	Valid clock mode	Default clock
DS3-E3 (Combo TDM OC48 FRU)	ANSI	DS3	DCR, ACR, INTERNAL	INTERNAL
		DS1	DCR, INTERNAL	INTERNAL
	ETSI	E3	DCR, ACR, INTERNAL	INTERNAL
T1 (Combo TDM OC48 FRU)	ANSI	DS1	DCR, INTERNAL, ACR	INTERNAL
E1 (Combo TDM OC48 FRU)	ETSI	E1	DCR, INTERNAL	INTERNAL

When a TDM SFP is interloped with a 3926 port, all timing modes that are SFP supports are supported except for EPAR.

Step	Action

Recovering the clock from OC3

1 Recover the clock from OC3:

```
port tdm oc3-stm1 set oc3-stm1 <tdmPortName> clock-mode
<internal | DCR-EPAR>
```

where

oc3-stm1 <tdmPortName>	is the TDM port name.
clock mode <internal | DCR-EPAR>	is the clock mode. The default is DCR-EPAR.

Recovering the clock from OC3-STM1 on a Combo TDM OC48 FRU module

2 Recover the clock from OC3-STM1 on a Combo TDM OC48 FRU module:

```
port tdm oc3-stm1-cem set oc3-stm1-cem <tdmPortName>
clock-mode <internal | DCR-EPAR>
```

where

oc3-stm1 <tdmPortName>	is the TDM port name.
clock mode <internal | DCR-EPAR>	is the clock mode. The default is internal.

Copyright© 2022 Ciena® Corporation

Recovering the clock from E1

3 Recover the clock from E1:

```
port tdm e1 set e1 <tdmPortName> clock-mode <internal |
DCR-EPAR | adaptive | recovered>
```

where

e1 <tdmPortName>	is the TDM port name.
clock mode <internal | DCR-EPAR | adaptive | recovered>	is the clock mode. The default is DCR-EPAR. ***Note:*** The “recovered” clock mode is not supported if the port is on a Combo TDM OC48 FRU module.

Recovering the clock from T1

4 Recover the clock from T1:

```
port tdm t1 set t1 <tdmPortName> clock-mode <internal |
DCR-EPAR | adaptive | recovered>
```

where

t1 <tdmPortName>	is the TDM port name.
clock mode <internal | DCR-EPAR>	is the clock mode. The default is DCR-EPAR. ***Note:*** The “recovered” clock mode is not supported if the port is on a Combo TDM OC48 FRU module.

Returning the clock to default values

5 Return the clock to default values:

```
port tdm <oc3-stm1|e1> unset <oc3-stm1|e1> <tdmPortName>
clock-mode
```

where

oc3-stm1 | e1	selects the TDM port to return to default.
<tdmPortName>	is the TDM port name.

Returning the clock to default values on a Combo TDM OC48 FRU module

6 Return the clock to default values on a Combo TDM OC48 FRU module:

```
port tdm <oc3-stm1-cem|t1> unset <oc3-stm1-cem|t1-cem>
<tdmPortName> clock-mode
```

where

oc3-stm1-cem | t1	selects the TDM port to return to default.
<tdmPortName>	is the TDM port name.

Copyright© 2022 Ciena® Corporation

Configuring the clock to adaptive clock-mode for Combo-FRU port T1

7 Setting the clock to adaptive clock-mode for Combo-FRU port T1:

```
port tdm t1 set port <tdmPortName> clock-mode adaptive
master-clock-src-tdm-vc <TDM VC Name>
```

where

<tdmPortName> is the TDM port name.

`<TDM VC Name>` is the TDM virtual circuit name.

Configuring the clock to adaptive clock-mode for Combo-FRU port DS3

8 Setting the clock to adaptive clock-mode for Combo-FRU port DS3:

```
port tdm ds3-e3 set ds3-e3 <tdmPortName> clock-mode
adaptive master-clock-src-tdm-vc <TDM VC Name>
```

where

<tdmPortName> is the TDM port name.

`<TDM VC Name>` is the TDM virtual circuit name.

Overwrite the clock from ports that do not support SSM

9 For ports which don't support SSM, overwrite the clock quality by applying forced-ql irrespective of the clock-mode:

```
sync tdm input create ref <Name> vc <VC Name> clock-mode
acr forced-ql prs
```

where

<Name> is the timing reference.

`<VC Name>` is the TDM virtual circuit name.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-22
Configuring TDM global synchronization

Configure TDM global synchronization:

- reversion-mode
- wait-to-restore

Step	Action

To configure reversion-mode

1 Configure reversion-mode:

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

where

reversion-mode <non-revertive|revertive>	Specifies whether the clock selection algorithm selects the highest available quality clock, even if the current clock is above the required quality level threshold. Default is non-revertive.

To configure wait-to-restore mode

2 Configure wait-to-restore mode:

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

where

wait-to-restore <MINUTES: 0..12>	Specifies the amount of time, in minutes (0 to 12) before reverting to the highest available quality clock port.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-23
Configuring TDM synchronization frequency protection groups

The TDM synchronization frequency protection group (SFPG) is assigned to a TDM 3926 port. A TDM SFPG is a list of clock input references. When a clock state changes, a clock selection algorithm determines the best clock source. If the current selected clock source is below the threshold quality level and the new clock is at a higher quality level, then the new clock is selected as the current clock source. When revertive mode is set, a higher quality level clock is always selected after the wait-to-restore time expires.

The input clock references can be line, internal, ACR or DCR.

The clock selection algorithm executes whenever one of these changes occur:

- The state of an input reference changes.
- An attribute of an input reference changes.
- An attribute of the TDM SFPG changes.

The clock selection algorithm is executed on each protection group as follows:

- Generates the list of good input references based on the clock state.
  - For line timing reference, this is the port state.
  - For ACR/DCR, this is the PWE state.
  - For internal, this is the clock quality.
- If more than one input reference is good, select the input reference with the highest override priority.
- If more than one input is still available, select the input reference with the highest forced quality level.
- If more than one input reference is still available, select the input reference with the highest priority.
- If more than one input reference is available, one is selected at random.
- The quality level of the input reference is also taken into consideration.

Copyright© 2022 Ciena® Corporation

Step	Action

To configure a TDM synchronization frequency protection group

1 Configure a TDM SFPG:

```
port tdm sync frequency protection-group create group
<group-name> threshold-ql <ql>
```

where

group <group-name>	Specifies the group name of up to 31 characters in length.
threshold-ql <ql>	Specifies the threshold quality level. The clock quality level is defined when a new clock input reference is selected. This occurs if the current selected clock falls below this quality level. The quality level has a different values based on the TDM mode: • ANSI values are — <dnu | prov | prs | smc | ssua | st2 | st3 | st3e | stu> Default is st3. • ETSI values are — <dus | prc | sec | ssub | tnc>, Default is sec. If the TDM mode is changed, then the threshold_ql changes to the default quality level for that TDM mode.

To add clock input references

2 Add clock input references:

```
port tdm sync frequency protection-group add group
<group-name> input-ref <input-ref-name>
```

where

input-ref <input-ref-name>	Specifies the TDM input clock timing reference.

To remove clock input references

3 Remove clock input references:

```
port tdm sync frequency protection-group remove group
<group-name> input-ref <input-ref-name>
```

where

input-ref <input-ref-name>	Specifies the TDM input clock timing reference.

Copyright© 2022 Ciena® Corporation

To remove clock input references

4 Remove clock input references:

```
port tdm sync frequency protection-group remove group
<group-name> input-ref <input-ref-name>
```

where

input-ref <input-ref-name>	Specifies the TDM input clock timing reference.

To delete a protection group

5 Delete a protection group:

```
port tdm sync frequency protection-group delete group
<group-name>
```

Note: A group cannot be deleted if it is still assigned to a TDM port.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-24
Clearing TDM port statistics

Port statistics displayed using the CLI show command are allowed to be cleared, but this does not affect PMs managed by performance management. Statistics may be buffered for up to 10 seconds, so it is possible to see a small jump in statistics after they are cleared when using the show command. For example, you may have just cleared port statistics, but the count when disabled may already show 10 errored seconds.

This procedure applies to ports on the TDM SFP FRU, to ports on the Combo TDM OC48 FRU module, and to ports on the 3926.

Step	Action
1	Clear TDM port statistics: `port tdm clear port <tdmPortName> statistics` where <tdmPortName> is the TDM port name. **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 10-25
Configuring TDM profiles

For every virtual circuit, the pseudowire type and payload size must be specified. For pseudowires using **most** ports on the 3926, the TDM profile contains those specifications.

Pseudowires using the ports on the Combo TDM OC48 FRU module are different because they do not require a TDM profile. In the absence of a TDM profile, the payload size and PW type are determined as follows:

- The payload size can be set by using the payload-size parameter. This is an optional parameter that applies only to virtual circuits using the Combo TDM OC48 FRU module. If the parameter is not set explicitly, a default value applies. For details, see "Configuring TDM virtual circuits".
- The PW type (CESoP or SAToP) is determined by the combination of the port and the access-circuit channelization, as listed in the table "TDM SFP port information".

For all pseudowires other than those that use the ports on the Combo TDM OC48 FRU module, you must configure a TDM profile before a virtual circuit can be configured.

Configure a TDM profile to contain the TDM parameters that are used by the TDM virtual circuit and the MPLS Layer 2 Virtual Private Network (L2VPN) pseudowire.

A TDM profile can be used by one or more TDM virtual circuits and MPLS L2VPN pseudowires. Note that once a profile has been associated with a virtual circuit or MPLS L2VPN pseudowire, the parameters in the profile cannot be changed without first disassociating the profile from all the virtual circuits or MPLS L2VPNs that it is associated with. Alternately, you can create another profile with different parameters and associate it with the required virtual circuits and MPLS L2VPN pseudowires.

> ***Note 1:*** A TDM profile cannot be deleted without first disassociating it from a TDM virtual circuit or an MPLS L2VPN. For the commands, refer to *39XX/51XX Service Delivery, Aggregation, and Virtualization Switches Command Reference*.

Copyright© 2022 Ciena® Corporation

Step	Action
1	Configure the TDM profile: `virtual-circuit tdm create tdm-profile <TDM Profile Name> pwType {satop|cesop} payload-size <bytes>` where tdm-profile <TDM Profile Name> — is the name of the TDM profile. pwType {satop|cesop} — is the pseudowire type, which is one of SAToP or CESoP.

Copyright© 2022 Ciena® Corporation

where

payload-size <bytes> ***Note:*** Payload size is measured as follows: • <bytes> for PWType=satop • <multiple of number of channels> for PWType=cesop	is the payload size. If the pwType is satop, payload-size is • for T1, multiples of 24 to a maximum of 192 bytes • for E1, multiples of 32 to a maximum of 256 bytes If the pwType is cesop, payload-size depends on the number of channels used by the attachment circuit that are referenced along with the TDM profile when the TDM VC is created. The rules for payload-size where pwType is cesop are: • Settable range is 1 to 192 for T1 and 1 to 256 for E1. • Must be a multiple of the number of channels configured on the attachment circuit that is used. For example, 10 channels use payload size of 10, 20, 30, and so on. • Payload-size maximum must not exceed 64 times the number of channels. For example, payload-size for two channels cannot be greater than 128. • Payload-size divided by the number-of-channels times 125 microseconds must not exceed configured PDV value of TDM-VC. If the payload-size is for a 3926 port, the settable range is 1 to 256. • for E1, the payload-size is 256 bytes • for T1, the payload-size is 192 bytes The CESoPSN PW default payload bytes are calculated as follows: • Payload Size (L) = * * Number of Time Slots (N) * Delay in Milliseconds (D) The default Milliseconds that must be supported are as follows: • 8 milliseconds (5 is recommended) when N is one, that is payload size is 64 bytes • 4 milliseconds when N is 2, 3 or 4, that is payload size is 32 * N • 1 milliseconds when N > 5, that is payload size is 8 * N For example, for 4xDS0, the payload size is 32 * 4 = 128 bytes.

Copyright© 2022 Ciena® Corporation

where

These defaults can be overridden by using the TDM payload bytes sub-tly interface parameter with Type set to 0x04, length set to 0x04 and value set to PSN MTU value as follows:

- For SAToP, PSN MTU value represents the number of payload bytes which does not include PW header, tunnel header and layer 2 header
- For CESoPSN, PSN MTU value represents “integer multiple of timeslots”. If MTU value is set to 10, it represents 10 * N timeslots. These rules/factors apply for payload size:
  - Payload size is an integer multiple of the number of timeslots as previously indicated.
  - Payload size is between 1 to 256 maximum
  - Payload size/time-slots cannot exceed 64.
  - Payload size/time-slots * 125 usec must be less than the user configured PDV.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-26
Configuring attachment circuits

Configure an attachment circuit to represent the client UNI port as a single logical entity.

For E1, channel 0 is always reserved. In PCM30, you must create a separate CESoP pseudowire for the channel 16 signaling channel only, and then bundle the other channels in other CESoP pseudowires.

You can

- create a SAToP attachment circuit
- create a CESoP attachment circuit
- add channels to a CESoP attachment-circuit

  Note: For attachment circuits attached to the ports on the Combo TDM OC48 FRU module, only single channels are supported, and only one channelization per port is supported.

- remove channels from a CESoP attachment-circuit
- set attributes for a CESoP attachment-circuit
- unset attributes for a CESoP attachment-circuit

  Note: An attachment circuit cannot be deleted if it is associated with a TDM virtual circuit.

Step	Action

To create an SAToP attachment circuit

1 Create an SAToP attachment circuit that uses the specified TDM port instance:

```
attachment-circuit satop create ac <ac-name> port <TDM
port> [channel <channel-num>][channel-type
<ds1|e1|none>] rtp <on|off>.
```

where

ac <ac-name>	is the attachment circuit.
port <TDM port>	is the TDM port.

Copyright© 2022 Ciena® Corporation

where

channel <channel-num>	is the channel set.
channel-type <ds1|e1|none>	is the channel type of ds1, e1 or none.
rtp <on|off>	specifies a fixed RTP header sent as part of the TDM PWE packet. The fixed RTP header as the ssrc id set to 0 and the payload type set to 0. Default is on. ***Note 1:*** RTP can be on or off for CSoP only. The RTP header must be on when the clocking-mode is DCR. For all other clocking modes, the RTP header may be on or off. Turning off the RTP header reduces the packet size. It is also intended for interworking with other third-party SFPs. ***Note 2:*** The RTP attribute also applies to 3926 ports. ***Note 3:*** Disable RTP on an OC48 Combo TDM FRU port if that port is interoperating with a one of the following transceivers: XCVR-CEP001, XCVR-CEP002 XCVR-CEP003.

To create a CESoP attachment circuit

2 Create a CESoP attachment circuit on the specified TDM port instance using the selected list of channels:

```
attachment-circuit cesop create ac <ac-name> port <tdm-
port-name> channel <channel-num> [channel-type <ds0-
e0|ds1-e1|sts1-vc3|sts3c-vc4c|stsl2-vc4-4c rtp <on|off>
```

where

ac <ac-name>>	is the name of the attachment circuit. The characters !%,?* are not permitted.
port <TDM port>	is the TDM port.

Copyright© 2022 Ciena® Corporation

where

channel <num>	is the channel set. If the TDM port type is T1, the channel set can be any unused channel in the range of 1 to 24. If the TDM port type is E1, the channel set can be any unused channel in the range of 1 to 31. For the E1 port type, channel 0 is reserved for signaling and cannot be used in a CESoP attachment circuit. If the TDM port type is OC3-STM1, OC12-STM4 or DS3-E3, refer to Table 10-1 on page 10-3 for allowed channels. If the TDM port type is OC3-STM1-CEM, OC12-STM4-CEM, OC48-STM16-CEM, or DS3-E3, refer to Table 10-1 on page 10-3 for allowed channels.
channel-type <ds0-e0|ds1-e1|sts1-vc3|sts3c-vc4c|stsl2-vc4-4c>	is the optional channel type. This attribute applies if the TDM port type is OC3-STM1, OC12-STM4, DS3-E3 or E1. Refer to Table 10-1 on page 10-3 for allowed channels.
rtp <on|off>	specifies a fixed RTP header sent as part of the TDM PWE packet. This attribute only applies to attachment circuits associated with TDM ports. The fixed RTP header as the ssrc id set to 0 and the payload type set to 0. Default is on. The RTP header is required when the TDM ports are connected to other ports that have their timing set to dcr-epar. This is necessary for interworking with SFP TDM ports that have their timing set to dcr-epar. ***Note 1:*** RTP can be on or off for CSoP only. The RTP header must be on when the clocking-mode is DCR. For all other clocking modes, the RTP header may be on or off. Turning off the RTP header reduces the packet size. It is also intended for interworking with other third-party SFPs. ***Note 2:*** The RTP attribute also applies to ports on the 3926, and to ports on the Combo TDM OC48 FRU module. ***Note 3:*** Disable RTP on an OC48 Combo TDM FRU port if that port is interoperating with a one of the following transceivers: XCVR-CEP001, XCVR-CEP002 XCVR-CEP003.

Copyright© 2022 Ciena® Corporation

To add channels to a CESoP attachment-circuit

3 Add channels to a CESoP attachment circuit:

```
attachment-circuit cesop add ac <ac-name> channel
<channel-list>
```

where

ac <ac-name>	is the attachment circuit.
channel <channel-list>	is the channel set. For the T1 port type, the channel set can be any unused channel in the range of 1 to 24. For the E1 port type, the channel set can be any unused channel in the range of 1 to 31. Channel 0 is reserved for signaling. For the OC3-STM1 port type, the channel can only be set to 0. ***Note:*** For ports on the Combo TDM OC48 FRU module, channels cannot be added.

To set CESoP attachment-circuit attributes

4 Set the attributes of the CESoP attachment circuit:

```
attachment-circuit cesop set ac <ac-name> [master-clock-
src <on|off>]
```

where

ac <ac-name>	is the attachment circuit.
[master-clock-src <on|off>]	indicates whether this circuit is the master clock source in adaptive clock mode. The default value is off. If changing from on to off, make sure the associated TDM VC is not currently used as a master clock source for one of the TDM ports that have its clock mode set to adaptive. If setting from off to on, make sure there is no other attachment circuit on the targeted port with its master-clock-src flag set to on. Only CESoP attachment circuits associated with TDM ports 1 to 8 can be designated as the master clock source. ***Note 1:*** The master clock source is not supported on TDM SFPs. ***Note 2:*** There can only be one CESoP attachment circuit for each TDM port designated as the master clock source. Only the TDM VC that is associated with a CESoP attachment circuit that has its master clock source flag set to on can be used as a master clock source TDM VC.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-27
Configuring signal degrade

Configure signal degrade and return signal degrade to default values.

Step	Action

To configure signal degrade

1 Configure signal degrade:

```
port tdm <tdmPortType> set <tdmPortType> signal-degrade
<NUMBER: 5..9>
```

where

<tdmPortType>	is the TDM port type. This value can be OC3/STM1, or OC12/STM4 SFPs For ports on the Combo TDM OC48 FRU module, this value can be oc12-stm4-cem or oc3-stm1-cem.
signal-degrade <NUMBER: 5..9>	is the bit error rate for signal degrade.

To return signal degrade to default values

2 Return signal degrade to default values:

```
port tdm <tdmPortType> unset <tdmPortType> signal-degrade
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-28
Configuring signal fail

Configure signal fail and return signal fail to default values.

Step	Action

To configure signal fail

1 Configure signal fail:

```
port tdm <tdmPortType> set <tdmPortType> signal-fail
<NUMBER: 3..5>
```

where

<tdmPortType>	is the TDM port type. This value can be OC3/STM1 OR OC12/STM4 SFPs. For ports on the Combo TDM OC48 FRU module, this value can be oc12-stm4-cem or oc3-stm1-cem.
signal-fail <NUMBER: 3..5>	is the bit error rate for signal fail.

To return signal degrade to default values

2 Return signal degrade to default values:

```
port tdm <tdmPortType> unset <tdmPortType> signal-fail
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-29
Configuring OC3 port loopbacks

Configure, enable and disable loopbacks on OC3 ports as required by the network plan.

Loopbacks can be set at the channel level or the port level. When loopbacks are configured at the channel level, only one channel can be looped back at any one time. If the loopback channel is also used by an attachment circuit, both the loopback and the attachment circuit must have the same channelization.

> ***Note:*** The “port” channelization type is not available for OC3 ports on the Combo TDM OC48 FRU module.

Step	Action

To set OC3 port loopback attributes

1 Set OC3 port loopback:

```
port tdm oc3-stm1 loopback set oc3-stm1 <oc3-stm1>
[channel <NUMBER: 1..83>] [type <port | sts1-vc3 | sts3c-
vc4c>] [mode <facility|terminal>]
```

where

oc3-stm1 <oc3-stm1>	is the OC3 port name.
channel <NUMBER: 0..83>	is the channel number. The default channel is 0.
type <sts1-vc3 | sts3c-vc4c>	is the loopback channelization type. The default value is port.
mode <facility | terminal>	is the loopback mode. Valid values are: • facility: indicates that the loopback is toward the system where the transmission to the SFP is received back to the system. • terminal: indicates that the loopback is toward the user where the receive from the SFP is transmitted back to the user. The default value is terminal.

Copyright© 2022 Ciena® Corporation

To set OC3 port attributes on a port on a Combo TDM OC48 FRU module

2 Set loopback on an OC3 port on a Combo TDM OC48 FRU module:

```
port tdm oc3-stm1-cem loopback set oc3-stm1-cem <oc3-stm1-cem> [channel <NUMBER: 0..83>] [type <sts1-vc3 | sts3c-vc4c>] [mode <facility|terminal>]
```

where

oc3-stm1-cem <oc3-stm1-cem>	is the OC3 port name.
channel <NUMBER: 0..83>	is the channel number. The default channel is 0.
type <sts1-vc3 | sts3c-vc4c | ds1>	is the loopback channelization type. ***Note:*** For loopback with DS1 channelization , a PW must be set before enabling the remote loopback.
mode <facility | terminal>	is the loopback mode. Valid value is: • facility: indicates that the loopback is toward the system where the transmission to the SFP is received back to the system. • terminal: indicates that the loopback is toward the user where the receive from the SFP is transmitted back to the user. The default value is terminal.

To enable OC3 port loopback

3 Enable OC3 port loopback:

```
port tdm oc3-stm1 loopback enable oc3-stm1 <oc3-stm1>
```

where

oc3-stml <oc3-stml >	is the OC3 port name.

To enable OC3 port loopback on a port on a Combo TDM OC48 FRU module

4 Enable loopback on an OC3 port on a Combo TDM OC48 FRU module:

```
port tdm oc3-stm1-cem loopback enable oc3-stm1-cem <oc3-stm1-cem>
```

where

oc3-stml-cem <oc3-stml -cem >	is the OC3 port name.

Copyright© 2022 Ciena® Corporation

To disable OC3 port loopback

5 Disable OC3 port loopback:

```
port tdm oc3-stm1 loopback disable oc3-stm1 <oc3-stm1>
```

where

oc3-stml <oc3-stml >	is the OC3 port name.

To disable OC3 port loopback on a port on a Combo TDM OC48 FRU module

6 Disable loopback on an OC3 port on a Combo TDM OC48 FRU module:

```
port tdm oc3-stm1-cem loopback disable oc3-stm1-cem <oc3-
stm1-cem>
```

where

oc3-stml-cem <oc3-stml -cem>	is the OC3 port name.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-30
Configuring OC12 port loopbacks

Configure, enable and disable loopbacks on OC12 ports as required by the network plan.

Loopbacks can be set at the channel level or the port level. When loopbacks are configured at the channel level, only one channel can be looped back at any one time. If the loopback channel is also used by an attachment circuit, both the loopback and the attachment circuit must have the same channelization.

> ***Note:*** The “port” channelization type is not available for OC12 ports on the Combo TDM OC48 FRU module.

Step	Action

To set OC12 port loopback attributes

1 Set OC12 port loopback:

```
port tdm oc12-stm4 loopback set oc12-stm4 <oc12-stm4>
[channel <NUMBER: 1-11>] [type <port | sts1-vc3 | sts12c-
vc4-4c | sts3cvc4c>] [mode <facility|terminal>]
```

where

oc12-stm4 <oc12-stm4>	is the OC12 port name.
channel <NUMBER: 0-11>	is the channel number. The default channel is 0.
type <port | sts1-vc3 | sts12c-vc4-4c | sts3cvc4c>	is the loopback channelization type. The default value is port.
mode <facility | terminal>	is the loopback mode. Valid values are: • facility: indicates that the loopback is toward the system where the transmission to the SFP is received back to the system. • terminal: indicates that the loopback is toward the user where the receive from the SFP is transmitted back to the user. The default value is terminal.

Copyright© 2022 Ciena® Corporation

To set OC12 port attributes on a port on a Combo TDM OC48 FRU module

2 Set loopback on an OC12 port on a Combo TDM OC48 FRU module:

```
port tdm oc12-stm4-cem loopback set oc12-stm4-cem <oc12-
stm4-cem> [channel <NUMBER: 0..335>] [type <sts1-vc3 |
sts12c-vc4-4c | sts3cvc4c>] [mode <facility|terminal>]
```

where

oc12-stm4-cem <oc12-stm4-cem>	is the OC12 port name.
channel <NUMBER: 0..335>	is the channel number. The default channel is **0.**
type <sts1-vc3 | sts12c-vc4-4c | sts3c-vc4c | ds1>	is the loopback channelization type. ***Note:*** For loopback with DS1 channelization , a PW must be set before enabling the remote loopback.
mode <facility | terminal>	is the loopback mode. Valid value is: • facility: indicates that the loopback is toward the system where the transmission to the SFP is received back to the system. • terminal: indicates that the loopback is toward the user where the receive from the SFP is transmitted back to the user. The default value is terminal.

To enable OC12 port loopback

3 Enable OC12 port loopback:

```
port tdm oc12-stm4 loopback enable oc12-stm4 <oc12-stm4>
```

where

oc12-stm4 <oc12-stm4>	is the OC12 port name.

To enable OC12 port loopback on a port on a Combo TDM OC48 FRU module

4 Enable loopback on an OC12 port on a Combo TDM OC48 FRU module:

```
port tdm oc12-stm4-cem loopback enable oc12-stm4-cem
<oc12-stm4-cem>
```

where

oc12-stm4-cem <oc12-stm4-cem>	is the OC12 port name.

Copyright© 2022 Ciena® Corporation

To disable OC12 port loopback

5 Disable OC12 port loopback:

```
port tdm oc12-stm4 loopback disable oc12-stm4 <oc12-stm4>
```

where

oc12-stm4 <oc12-stm4 >	is the OC12 port name.

To disable OC12 port loopback on a port on a Combo TDM OC48 FRU module

6 Disable loopback on an OC12 port on a Combo TDM OC48 FRU module:

```
port tdm oc12-stm4-cem loopback enable oc12-stm4-cem
<oc12-stm4-cem>
```

where

oc12-stm4-cem <oc12-stm4-cem>	is the OC12 port name.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-31
Configuring OC48 port loopbacks

Configure, enable and disable loopbacks on OC48 ports as required by the network plan.

The OC48 port is on the Combo TDM OC48 FRU module. On this port, loopbacks can be set at the channel level only. Only one channel can be looped back at any one time. If the loopback channel is also used by an attachment circuit, both the loopback and the attachment circuit must have the same channelization.

> ***Note:*** The “port” channelization type is not available for the OC48 port on the Combo TDM OC48 FRU module.

Step	Action

To set OC48 port attributes on a port on a Combo TDM OC48 FRU module

1 Set loopback on an OC48 port on a Combo TDM OC48 FRU module:

```
port tdm oc48-stm16-cem loopback set oc48-stm16-cem
<oc48-stm16-cem> [channel <NUMBER: 0..1343>] [type <sts1-
vc3 | sts12c-vc4-4c | sts3cvc4c>] [mode
<facility|terminal>]
```

where

oc48-stm16-cem <oc48-stm16-cem>	is the OC48 port name.

Copyright© 2022 Ciena® Corporation

where

channel <NUMBER: 0..1343>	is the channel number. The default channel is 0.
type <sts1-vc3 | sts12c-vc4-4c | sts3c-vc4c | sts48-vc4-16c | ds1>	is the loopback channelization type. ***Note:*** For loopback with DS1 channelization , a PW must be set before enabling the remote loopback.
mode <facility | terminal>	is the loopback mode. Valid value is: • facility: indicates that the loopback is toward the system where the transmission to the SFP is received back to the system. • terminal: indicates that the loopback is toward the user where the receive from the SFP is transmitted back to the user. The default value is terminal.

To enable OC48 port loopback on a port on a Combo TDM OC48 FRU module

2 Enable loopback on an OC48 port on a Combo TDM OC48 FRU module:

```
port tdm oc48-stm16-cem loopback enable oc48-stm16-cem
<oc48-stm16-cem>
```

where

oc48-stm16-cem <oc48-stm16-cem>	is the OC48 port name.

To disable OC48 port loopback on a port on a Combo TDM OC48 FRU module

3 Disable loopback on an OC48 port on a Combo TDM OC48 FRU module:

```
port tdm oc12-stm4-cem loopback enable oc12-stm4-cem
<oc12-stm4-cem>
```

where

oc48-stm16-cem <oc48-stm16-cem>	is the OC48 port name.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-32
Configuring DS3 port loopbacks

Configure, enable and disable loopbacks on DS3 ports as required by the network plan.

The DS3 ports on the Combo TDM OC48 FRU module support loopback with DS1 channelization and port loopback. For loopback with DS1 channelization, a PW must be set before enabling the remote loopback.

Step	Action

To set DS3 port loopback attributes

1 Set DS3 port loopback:

```
port tdm ds3-e3 loopback set ds3-e3 <ds3-e3> [mode
<facility | terminal>]
```

where

ds3-e3 <ds3-e3>	is the port name.
mode <facility | terminal>	is the loopback mode. Valid values are: • facility: indicates that the loopback is toward the system where the transmission to the SFP is received back to the system. • terminal: indicates that the loopback is toward the user where the receive from the SFP is transmitted back to the user. The default value is terminal.

To enable DS3 port loopback

2 Enable DS3 port loopback:

```
port tdm ds3-e3 loopback enable ds3-e3 <ds3-e3>
```

where

ds3-e3 <ds3-e3>	is the DS3 port name.

To disable DS3 port loopback

3 Disable DS3 port loopback:

```
port tdm ds3-e3 loopback disable ds3-e3 <ds3-e3>
```

where

ds3-e3 <ds3-e3 >	is the DS3 port name.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-33
Configuring E1 port loopbacks

Configure, enable and disable loopbacks on E1 ports as required by the network plan.

Step	Action

To set E1 port loopback attributes

1 Set E1 port loopback:

```
port tdm e1 loopback set e1 <e1> [mode <facility |
terminal>]
```

where

e1 <e1>	is the port name.
mode <facility | terminal>	is the loopback mode. Valid values are: • facility: indicates that the loopback is toward the system where the transmission to the SFP is received back to the system. • terminal: indicates that the loopback is toward the user where the receive from the SFP is transmitted back to the user. The default value is terminal.

To enable E1 port loopback

2 Enable E1 port loopback:

```
port tdm e1 loopback enable e1 <e1>
```

where

e1 <e1>	is the E1 port name.

To disable E1 port loopback

3 Disable E1 port loopback:

```
port tdm e1 loopback disable e1 <e1>
```

where

e1 <e1 >	is the E1 port name.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-34
Configuring T1 port loopbacks

Configure, enable and disable loopbacks on T1 ports as required by the network plan.

> ***Note:*** Loopcodes cannot be enabled or transmitted at the same time loopback is activated.

For T1 ports on the Combo TDM OC48 FRU module, only the port loopback is supported.

Step	Action

To set T1 port loopback attributes

1 Set T1 port loopback:

```
port tdm t1 loopback set t1 <t1> [mode <facility |
terminal>]
```

where

t1 <t1>	is the t1 port name.
mode <facility | terminal>	is the loopback mode. Valid values are: • facility: indicates that the loopback is towards the TDM client side, where TDM stream received on T1 interface is looped back to the TDM client. • terminal: indicates that the loopback is towards the PSN side where TDM stream recovered from IWF is looped back to the remote node. The default value is terminal.

Copyright© 2022 Ciena® Corporation

To set T1 port attributes on a port on a Combo TDM OC48 FRU module

2 Set loopback on a T1 port on a Combo TDM OC48 FRU module:

```
port tdm t1 loopback set t1 <t1> [mode
<facility|terminal>]
```

where

t1 <t1>	is the t1 port name.
mode <facility | terminal>	is the loopback mode. Valid value is: • facility: indicates that the loopback is towards the TDM client side, where TDM stream received on T1 interface is looped back to the TDM client. • terminal: indicates that the loopback is towards the PSN side where TDM stream recovered from IWF is looped back to the remote node. The default value is terminal.

To enable T1 port loopback

3 Enable T1 port loopback:

```
port tdm t1 loopback enable t1 <t1>
```

where

t1 <t1>	is the T1 port name.

To disable T1 port loopback

4 Disable T1 port loopback:

```
port tdm t1 loopback disable t1 <t1>
```

where

t1 <t1 >	is theT1 port name.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-35
Configuring TDM virtual circuits

Configure a virtual circuit to represent the pseudowire-specific and PSN-specific processing that the PWE functionality performs. The virtual circuit is used for proxy PSN-specific processing.

Configuring a virtual circuit requires:

- a corresponding attachment circuit
- For the procedure to create a TDM profile, refer to "Configuring TDM profiles" on page 10-96.

Note: To change the attachment circuit associated with a virtual circuit, delete and re-add the TDM virtual circuit with the new attachment circuit association.

Step	Action

To create a TDM virtual circuit of type MEF8

1 Create a TDM virtual circuit of type MEF8:

```
virtual-circuit tdm-mef8 create vc <vc> {ac <Attachment
Circuit Name>} [c-vid <NUMBER: 1-4094>] [c-pcp <NUMBER:
0-7>] {peer-mac <MAC address: XX:XX:XX:XX:XX:XX>} [pdv
<NUMBER: 1000-32000>] {in-ecn <NUMBER: 1-65535>} {out-ecn
<NUMBER: 1-65535>} {c-tpid <8100|9100|88a8>} {tdm-profile
<TDM Profile Name>} {cas <on | off>}
```

where

vc <vc>	is the virtual circuit name. The name of a TDM virtual circuit must be unique across all TDM virtual circuits.
ac <Attachment Circuit Name>	is the attachment circuit.
c-vid <NUMBER: 1-4094>	is the CVLAN tag to be inserted in the Ethernet frame. If a C-VID is used on one virtual circuit of the TDM SFP then all other virtual circuits for the same TDM SFP must also use a C-VID.
c-pcp <NUMBER: 0-7>	is the P-bits (priority) value. The default value is 7.
peer-mac <MAC address: XX:XX:XX:XX:XX:XX>	is the 6-byte unicast MAC address of the remote end of the PWE service.

Copyright© 2022 Ciena® Corporation

where

pdv <NUMBER: 1000-32000>

is the Packet Delay Variation (PDV) in microseconds. The range is 1000 to 32000, that is, 1 to 32 milliseconds. The pdv parameter determines the size of the jitter buffer used to store incoming TDM service packets. The value of the pdv parameter must be greater than Payload-size divided by the number of channels times 125 microseconds for the CESoP TDM service.

For SAToP services, the default value is 3000 microseconds.

For CESoP services, the default value is calculated based on the number of channels used according to this equation: ((payload-size modulo num-channels) + n) so that this sum * 125 >= 1000.

For interworking PWE module ports with OC3, OC12, STM-1, STM-4 and DS3 TDM SFPs, the pdv must be set to 32000 on the virtual circuit associated with the PWE module port. For interworking PWE module ports with E1 TDM SFPs, the pdv must be set to 16000 on the virtual circuit associated with the PWE module port. (The PWE module was available in earlier releases.) (The PWE module was available in earlier releases.)

This attribute is used for SFP TDM ports. For CSoP, TPoP and 3926 ports, if PDV is set on a non-supported SFP, a warning is issued and the value is ignored. The defaults are as follows:

- CSoP is 16ms
- TPoP is 32ms
- 3926 port is 3ms for SAToP
- 3926, CESoP uses this formula: [(payload/number of timeslots) +1+n] x125 >1000 where n is in the lowest integer in the range of 0-8 that satisfies the equation.

For TPoP using a TDM4GE SFP connected to a faceplate port on a 3926, using DS1 channelization, the minimum PDV is 4000, the maximum is 32000, and the default is 32000.

On the SFP, the maximum PDV can be reached under specific circumstances. The actual PDV is about 4 ms below the maximum value. For a CSoP, the realistic achievable PDV is 12 ms, and for TPoP it is 28 ms. If the PDV is set to 10 ms, the actual realistic achievable PDV is 6 ms. The PDV is specified in terms of microseconds, from 1000 up to 32000. The TDM SFPs and module ports only accept PDV in increments of ms, so all values are rounded up to the nearest ms. For example a PDV of 11111 microseconds is converted to 12 ms.

Copyright© 2022 Ciena® Corporation

where

pdv <NUMBER: 1000-32000>	For the ports on the Combo TDM OC48 FRU module, the default value, is not user-configurable.

(continued)

in-ecn <NUMBER: 1-65535>	is the incoming or local Ethernet Connection Identifier (EC-ID). On packets received from the PSN, the EC-ID of the ingress packet is used to find the local TDM service this packet belongs to. The value is in the range of 1 to 65535. It must be unique among all tdm-mef8 (in-ecn).
out-ecn <NUMBER: 1-65535>	is the outgoing or remote EC-ID that is inserted on the TDM packet sent out to the PSN, that is, the EC-ID of the remote peer. Set the out-ecn parameter to the in-ecn value at the remote peer. Note that there are no local checks to verify that the values correspond.
c-tpid <8100|9100|88a8>	is the customer TPID on the TDM virtual circuit. ***Note:*** Setting this attribute is the only way to control c-tpid for virtual circuits on ports on the Combo TDM OC48 FRU module.

Copyright© 2022 Ciena® Corporation

where

tdm-profile <TDM Profile Name>	is the TDM profile name. ***Note 1:*** This attribute is not used for SFP TDM ports. ***Note 2:*** This attribute is used for 3926 ports. ***Note 3:*** This attribute is not required virtual circuits on the ports of the Combo TDM OC48 FRU module.
cas <on | off>	indicates whether channel associated signaling is on or off. ***Note:*** This attribute is not supported for virtual circuits on ports on the Combo TDM OC48 FRU module.
payload-size <NUMBER>	indicates the payload size in bytes. This parameter is used only for virtual circuits on the ports of the Combo TDM OC48 FRU module. It is an optional parameter. If not explicitly set the default value, and the minimum and maximum allowable values depend on the channelization, as follows: • With DS1 channelization, the minimum payload size is 96, the maximum payload size is 192, and the default payload size is 192. • With E1 channelization, the minimum payload size is 128, the maximum payload size is 256, and the default payload size is 256. • With DS3 and DS3->STS-1 channelization, the minimum payload size is 256, the maximum payload size is 1500, and the default payload size is 512. • With E3 channelization, the minimum payload size is 256, the maximum payload size is 2000, and the default payload size is 512. • With STS-1 channelization, the minimum payload size is 256, the maximum payload size is 1980, and the default payload size is 512. • With STS-3 or STM1 channelization, the minimum payload size is 783, the maximum payload size is 3000, and the default payload size is 783. • With STS-12 or STM4 channelization, the minimum payload size is 783, the maximum payload size is 3100, and the default payload size is 783. • With STS-48 or STM16 channelization, the minimum payload size is 783, the maximum payload size is 4000, and the default payload size is 783. ***Note:*** Set the payload size to 783 for an OC48 Combo TDM FRU port if that port is interoperating with a one of the following transceivers: XCVR-CEP001, XCVR-CEP002 XCVR-CEP003.

Copyright© 2022 Ciena® Corporation

To modify parameters for a TDM virtual circuit of type MEF8

2 Modify parameters for a TDM virtual circuit of type MEF8:

```
virtual-circuit tdm-mef8 set vc <vc> [ac <Attachment
Circuit Name>] [c-vid <vid>] [c-pcp <pcp-value>] [peer-
mac <mac>] [pdv <1000-32000>] [in-ecn <1-65535>] [out-ecn
<1-65535>] [c-tpid <8100|9100|88a8>] [tdm-profile <TDM
Profile Name>] [cas <on | off>]
```

where

vc <vc>	is the virtual circuit.
ac <Attachment Circuit Name>	is the attachment circuit.
c-vid <NUMBER: 1-4094>	is the CVLAN tag to be inserted in the Ethernet frame. If a C-VID is used on one virtual circuit of the TDM SFP then all other virtual circuits for the same TDM SFP must also use a C-VID.
c-pcp <NUMBER: 0-7>	is the P-bits (priority) value. The default value is 7.
peer-mac <MAC address: XX:XX:XX:XX:XX:XX>	is the 6-byte unicast MAC address of the remote end of the PWE service.

Copyright© 2022 Ciena® Corporation

where

pdv <NUMBER: 1000-32000>

is the Packet Delay Variation (PDV) in microseconds. The range is 1000 to 32000, that is, 1 to 32 milliseconds. The pdv parameter determines the size of the jitter buffer used to store incoming TDM service packets. The value of the pdv parameter must be greater than Payload-size divided by the number of channels times 125 microseconds for the CESoP TDM service.

For SAToP services, the default value is 3000 microseconds.

For CESoP services, the default value is calculated based on the number of channels used according to this equation: ((payload-size modulo num-channels) + n) so that this sum * 125 >= 1000.

For interworking PWE module ports with OC3, OC12, STM-1, STM-4 and DS3 TDM SFPs, the pdv must be set to 32000 on the virtual circuit associated with the PWE module port. For interworking PWE module ports with E1 TDM SFPs, the pdv must be set to 16000 on the virtual circuit associated with the PWE module port. (The PWE module was available in earlier releases.)

This attribute is used for SFP TDM ports. For CSoP, TPoP and 3926 ports, if PDV is set on a non-supported SFP, a warning is issued and the value is ignored. The defaults are as follows:

- CSoP is 16ms
- TPoP is 32ms
- 3926 port is 3ms for SAToP
- 3926 port, CESoP uses this formula:: [(payload/number of timeslots) +1+n] x125 >1000 where n is in the lowest integer in the range of 0-8 that satisfies the equation.

For TPoP using a TDM4GE SFP connected to a faceplate port on a 3926, using DS1 channelization, the minimum PDV is 4000, the maximum is 32000, and the default is 32000.

On the SFP, the maximum PDV can be reached under specific circumstances. The actual PDV is about 4 ms below the maximum value. For a CSoP, the realistic achievable PDV is 12 ms, and for TPoP it is 28 ms. If the PDV is set to 10 ms, the actual realistic achievable PDV is 6 ms. The PDV is specified in terms of microseconds, from 1000 up to 32000. The TDM SFPs and module ports only accept PDV in increments of ms, so all values are rounded up to the nearest ms. For example a PDV of 11111 microseconds is converted to 12 ms.

Copyright© 2022 Ciena® Corporation

where

pdv <NUMBER: 1000-32000>	For the ports on the Combo TDM OC48 FRU module, the default value, is not user-configurable.

(continued)

in-ecn <NUMBER: 1-65535>	is the incoming or local Ethernet Connection Identifier (EC-ID). On packets received from the PSN, the EC-ID of the ingress packet is used to find the local TDM service this packet belongs to. The value is in the range of 1 to 65535. It must be unique among all tdm-mef8 (in-ecn).
out-ecn <NUMBER: 1-65535>	is the outgoing or remote EC-ID that is inserted on the TDM packet sent out to the PSN, that is, the EC-ID of the remote peer. Set the out-ecn parameter to the in-ecn value at the remote peer. Note that there are no local checks to verify that the values correspond.
c-tpid <8100|9100|88a8>	is the customer TPID on the TDM virtual circuit. ***Note:*** Setting this attribute is the only way to control c-tpid for virtual circuits on ports on the Combo TDM OC48 FRU module.

Copyright© 2022 Ciena® Corporation

where

tdm-profile <TDM Profile Name>	is the TDM profile name. ***Note 1:*** This attribute is not used for SFP TDM ports. ***Note 2:*** This attribute is used for 3926 ports. ***Note 3:*** This attribute is not required virtual circuits on the ports of the Combo TDM OC48 FRU module.
cas <on | off>	indicates whether channel associated signaling is on or off. ***Note:*** This attribute is not supported for virtual circuits on ports on the Combo TDM OC48 FRU module.
payload-size <NUMBER>	indicates the payload size in bytes. This parameter is used only for virtual circuits on the ports of the Combo TDM OC48 FRU module. It is an optional parameter. If not explicitly set the default value, and the minimum and maximum allowable values depend on the channelization, as follows: • With DS1 channelization, the minimum payload size is 96, the maximum payload size is 192, and the default payload size is 192. • With E1 channelization, the minimum payload size is 128, the maximum payload size is 256, and the default payload size is 256. • With DS3 and DS3->STS-1 channelization, the minimum payload size is 256, the maximum payload size is 1500, and the default payload size is 512. • With E3 channelization, the minimum payload size is 256, the maximum payload size is 2000, and the default payload size is 512. • With STS-1 channelization, the minimum payload size is 256, the maximum payload size is 1980, and the default payload size is 512. • With STS-3 or STM1 channelization, the minimum payload size is 783, the maximum payload size is 3000, and the default payload size is 783. • With STS-12 or STM4 channelization, the minimum payload size is 783, the maximum payload size is 3100, and the default payload size is 783. • With STS-48 or STM16 channelization, the minimum payload size is 783, the maximum payload size is 4000, and the default payload size is 783. ***Note:*** Set the payload size to 783 for an OC48 Combo TDM FRU port if that port is interoperating with a one of the following transceivers: XCVR-CEP001, XCVR-CEP002 XCVR-CEP003.

Copyright© 2022 Ciena® Corporation

To create a TDM virtual circuit of type Dry Martini

3 Create a TDM virtual circuit of type Dry Martini:

```
virtual-circuit tdm-dry-martini create vc <vc> {ac
<Attachment Circuit Name>} [c-vid <NUMBER: 1-4094>] [c-
pcp <NUMBER: 0-7>] {peer-mac <MAC address:
XX:XX:XX:XX:XX:XX>} [pdv <NUMBER: 3000-32000>] {in-label
<NUMBER: 1-1048575>} {out-label <NUMBER: 1-1048575>}
{tdm-profile <TDM Profile Name>} {cas <on | off>}
```

where

vc <vc>	is the virtual circuit name. The name of a TDM virtual circuit must be unique across all TDM virtual circuits.
ac <Attachment Circuit Name>	is the attachment circuit.
c-vid <NUMBER: 1-4094>	is the CVLAN tag to be inserted in the Ethernet frame. If a C-VID is used on one virtual circuit of the TDM SFP then all other virtual circuits for the same TDM SFP must also use a C-VID. The same restrictions apply to 3926 ports.
c-pcp <NUMBER: 0-7>	is the P-bits (priority) value. The default value is 7. The same restrictions apply to 3926 ports.
peer-mac <MAC address: XX:XX:XX:XX:XX:XX>	is the 6-byte unicast MAC address of the remote end of the PWE service.
pdv <NUMBER: 1000-32000>	is the Packet Delay Variation (PDV) in microseconds. The range is 1000 to 32000, that is, 1 to 32 milliseconds. The pdv parameter determines the size of the jitter buffer used to store incoming TDM service packets. The value of the pdv parameter must be greater than Payload-size divided by the number of channels times 125 microseconds for the CESoP TDM service. For SAToP services, the default value is 3000 microseconds. For CESoP services, the default value is calculated based on the number of channels used according to this equation: ((payload-size modulo num-channels) + n) so that this sum * 125 >= 1000. For interworking PWE module ports with OC3, OC12, STM-1, STM-4 and DS3 TDM SFPs, the pdv must be set to 32000 on the virtual circuit associated with the PWE module port. For interworking PWE module ports with E1 TDM SFPs, the pdv must be set to 16000 on the virtual circuit associated with the PWE module port. (The PWE module was available in earlier releases.) ***Note:*** PDV is not available on TDM SFPs.

Copyright© 2022 Ciena® Corporation

where

in-label <NUMBER: 1-1048575>	is the incoming or local, that is pseudowire, label. On packet receive from the PSN, the pseudowire label from the packet is used to find local TDM service. It must be unique among all tdm-dry-martini (in-label) values. The label must be in the configured static MPLS VC label range.
out-label <NUMBER: 1-1048575>	is the outgoing or remote label that is inserted on the TDM packet sent out to the PSN, that is, the pseudowire label of the remote peer. Set the value of the out-label parameter to the value of the in-label of the remote peer. The label must be in the configured static MPLS virtual circuit label range and available, that is, not already used by static MPLS virtual circuit or another TDM Dry Martini pseudowire. Note that there are no local checks to verify that the values correspond.
tdm-profile <TDM Profile Name>	is the TDM profile name.
cas <on | off>	indicates whether channel associated signaling is on or off.

To modify the parameters of a TDM virtual circuit of type Dry Martini

4 Modify the parameters of a TDM virtual circuit of type Dry Martini:

Note: This operation is only permitted if the TDM virtual circuit is not cross-connected.

```
virtual-circuit tdm-dry-martini set vc <vc> [ac
<Attachment Circuit Name>] [c-vid <NUMBER: 1-4094>] [c-
pcp <NUMBER: 0-7>] [peer-mac <MAC address:
XX:XX:XX:XX:XX:XX>] [pdv <NUMBER: 1000-32000>] [in-label
<NUMBER: 1-1048575>] [out-label <NUMBER: 1-1048575>]
[tdm-profile <TDM Profile Name>] [cas <on | off>]
```

where

vc <vc>	is the virtual circuit.
[ac <Attachment Circuit Name>]	is the attachment circuit.
[c-vid <NUMBER: 1-4094>]	is the CVLAN tag to be inserted in the Ethernet frame. If a C-VID is used on one virtual circuit of the TDM SFP then all other virtual circuits for the same TDM SFP must also use a C-VID.
[c-pcp <NUMBER: 0-7>]	is the P-bits (priority) value. The default value is 7.
[peer-mac <MAC address: XX:XX:XX:XX:XX:XX>]	is the 6-byte unicast MAC address of the Pseudowire at the remote end of the PWE service.

Copyright© 2022 Ciena® Corporation

where

[pdv <NUMBER: 1000-32000>] is the Packet Delay Variation (PDV) in microseconds. The range is 1000 to 32000, that is, 1 to 32 milliseconds. The pdv parameter determines the size of the jitter buffer used to store incoming TDM service packets. The value of the pdv parameter must be greater than Payload-size divided by the number of channels times 125 microseconds for the CESoP TDM service.

For SAToP services, the default value is 3000 microseconds.

For CESoP services, the default value is calculated based on the number of channels used according to this equation: ((payload-size modulo num-channels) + n) so that this sum * 125 >= 1000.

For interworking PWE module ports with OC3, OC12, STM-1, STM-4 and DS3 TDM SFPs, the pdv must be set to 32000 on the virtual circuit associated with the PWE module port. For interworking PWE module ports with E1 TDM SFPs, the pdv must be set to 16000 on the virtual circuit associated with the PWE module port. (The PWE module was available in earlier releases.)

This attribute is used for SFP TDM ports. For CSoP, TPoP and 3926 ports, if PDV is set on a non-supported SFP, a warning is issued and the value is ignored. The defaults are as follows:

- CSoP is 16ms
- TPoP is 32ms
- 3926 port is 3ms for SAToP
- 3926 port, CESoP uses this formula:
  - [(payload/number of timeslots) +1+n] x125 >1000 where n is in the lowest integer in the range of 0-8 that satisfies the equation.

On the SFP, the maximum PDV can be reached under specific circumstances. The actual PDV is about 4 ms below the maximum value. For a CSoP, the realistic achievable PDV is 12 ms, and for TPoP it is 28 ms. If the PDV is set to 10 ms, the actual realistic achievable PDV is 6 ms. The PDV is specified in terms of microseconds, from 1000 up to 32000. The TDM SFPs and module ports only accept PDV in increments of ms, so all values are rounded up to the nearest ms. For example a PDV of 11111 microseconds is converted to 12 ms.

Copyright© 2022 Ciena® Corporation

where

{in-label <NUMBER: 1-1048575>}	is the incoming or local, that is pseudowire, label. On packet receive from the PSN, the pseudowire label from the packet is used to find local TDM service. It must be unique among all tdm-dry-martini (in-label) values. The label must be in the configured static MPIS VC label range.
{out-label <NUMBER: 1-1048575>}	is the outgoing or remote label that is inserted on the TDM packet sent out to the PSN, that is, the pseudowire label of the remote peer. Set the value of the out-label parameter to the value of the in-label at the remote peer. The label must be in the configured static MPLS virtual circuit label range and available, that is, not already used by static MPLS virtual circuit or another TDM Dry Martini pseudowire. Note that there are no local checks to verify that the values correspond.
[tdm-profile <TDM Profile Name>]	is the TDM profile name.
cas <on | off>	indicates whether channel associated signaling is on or off.

Copyright© 2022 Ciena® Corporation

To create a TDM virtual circuit of type MPLS

5 Create a TDM virtual circuit of type MPLS:

Note: The peer mac address, in-label, and out-label are set in the MPLS L2VPN virtual circuit configuration.

```
virtual-circuit tdm-mpls create vc <vc-name> ac <ac-name>
[pdv <3000-32000> tdm-profile <TDM Profile Name> {cas <on
| off>}]
```

where

vc <vc-name>	is the virtual circuit.
ac <ac-name>	is the attachment circuit.
pdv <3000-32000>	s the Packet Delay Variation (PDV) in microseconds. The range is 1000 to 32000, that is, 1 to 32 milliseconds. The pdv parameter determines the size of the jitter buffer used to store incoming TDM service packets. The value of the pdv parameter must be greater than Payload-size divided by the number of channels times 125 microseconds for the CESoP TDM service. For SAToP services, the default value is 3000 microseconds. For CESoP services, the default value is calculated based on the number of channels used according to this equation: ((payload-size modulo num-channels) + n) so that this sum * 125 >= 1000. For interworking PWE module ports with OC3, OC12, STM-1, STM-4 and DS3 TDM SFPs, the pdv must be set to 32000 on the virtual circuit associated with the PWE module port. For interworking PWE module ports with E1 TDM SFPs, the pdv must be set to 16000 on the virtual circuit associated with the PWE module port. (The PWE module was available in earlier releases.) This attribute is used for SFP TDM ports. For CSoP, TPoP and 3926 ports, if PDV is set on a non-supported SFP, a warning is issued and the value is ignored. The defaults are as follows: • CSoP is 16ms • TPoP is 32ms • 3926 port is 3ms for SAToP • 3926 port, CESoP uses this formula:: [(payload/number of timeslots) +1+n] x125 >1000 where n is in the lowest integer in the range of 0-8 that satisfies the equation. For TPoP using a TDM4GE SFP connected to a faceplate port on a 3926, using DS1 channelization, the minimum PDV is 4000, the maximum is 32000, and the default is 32000.

Copyright© 2022 Ciena® Corporation

where

pdv <3000-32000> (continued)	On the SFP, the maximum PDV can be reached under specific circumstances. The actual PDV is about 4 ms below the maximum value. For a CSoP, the realistic achievable PDV is 12 ms, and for TPoP it is 28 ms. If the PDV is set to 10 ms, the actual realistic achievable PDV is 6 ms. The PDV is specified in terms of microseconds, from 1000 up to 32000. The TDM SFPs and module ports only accept PDV in increments of ms, so all values are rounded up to the nearest ms. For example a PDV of 11111 microseconds is converted to 12 ms. For the ports on the Combo TDM OC48 FRU module, the default value, is not user-configurable.

Copyright© 2022 Ciena® Corporation

where

tdm-profile <TDM Profile Name>	is the TDM profile name. ***Note:*** This attribute is not required virtual circuits on the ports of the Combo TDM OC48 FRU module.
cas <on | off>	indicates whether channel associated signaling is on or off. ***Note:*** This attribute is not supported for virtual circuits on ports on the Combo TDM OC48 FRU module.
payload-size <NUMBER>	indicates the payload size in bytes. This parameter is used only for virtual circuits on the ports of the Combo TDM OC48 FRU module. It is an optional parameter. If not explicitly set the default value, and the minimum and maximum allowable values depend on the channelization, as follows: • With DS1 channelization, the minimum payload size is 96, the maximum payload size is 192, and the default payload size is 192. • With E1 channelization, the minimum payload size is 128, the maximum payload size is 256, and the default payload size is 256. • With DS3 and DS3->STS-1 channelization, the minimum payload size is 256, the maximum payload size is 1500, and the default payload size is 512. • With E3 channelization, the minimum payload size is 256, the maximum payload size is 2000, and the default payload size is 512. • With STS-1 channelization, the minimum payload size is 256, the maximum payload size is 1980, and the default payload size is 512. • With STS-3 or STM1 channelization, the minimum payload size is 783, the maximum payload size is 3000, and the default payload size is 783. • With STS-12 or STM4 channelization, the minimum payload size is 783, the maximum payload size is 3100, and the default payload size is 783. • With STS-48 or STM16 channelization, the minimum payload size is 783, the maximum payload size is 4000, and the default payload size is 783. ***Note:*** Set the payload size to 783 for an OC48 Combo TDM FRU port if that port is interoperating with a one of the following transceivers: XCVR-CEP001, XCVR-CEP002 XCVR-CEP003.

Copyright© 2022 Ciena® Corporation

To modify the parameters of a TDM virtual circuit of type MPLS

6 Modify the parameters of a TDM virtual circuit of type MPLS:

Note: This operation is only permitted if the TDM virtual circuit is not cross-connected.

```
virtual-circuit tdm-mpls set vc <vc-name> [ac <ac-name>]
[pdv <3000-32000>][tdm-profile <TDM Profile Name>] [cas
<on | off>]
```

where

vc <vc-name>	is the virtual circuit.
ac <ac-name>	is the attachment circuit.
pdv <3000-32000>	s the Packet Delay Variation (PDV) in microseconds. The range is 1000 to 32000, that is, 1 to 32 milliseconds. The pdv parameter determines the size of the jitter buffer used to store incoming TDM service packets. The value of the pdv parameter must be greater than Payload-size divided by the number of channels times 125 microseconds for the CESoP TDM service. For SAToP services, the default value is 3000 microseconds. For CESoP services, the default value is calculated based on the number of channels used according to this equation: ((payload-size modulo num-channels) + n) so that this sum * 125 >= 1000. For interworking PWE module ports with OC3, OC12, STM-1, STM-4 and DS3 TDM SFPs, the pdv must be set to 32000 on the virtual circuit associated with the PWE module port. For interworking PWE module ports with E1 TDM SFPs, the pdv must be set to 16000 on the virtual circuit associated with the PWE module port. (The PWE module was available in earlier releases.) This attribute is used for SFP TDM ports. For CSoP, TPoP and 3926 ports, if PDV is set on a non-supported SFP, a warning is issued and the value is ignored. The defaults are as follows: • CSoP is 16ms • TPoP is 32ms • 3926 port is 3ms for SAToP • 3926 port, CESoP uses this formula:: [(payload/number of timeslots) +1+n] x125 >1000 where n is in the lowest integer in the range of 0-8 that satisfies the equation. For TPoP using a TDM4GE SFP connected to a faceplate port on a 3926, using DS1 channelization, the minimum PDV is 4000, the maximum is 32000, and the default is 32000.

Copyright© 2022 Ciena® Corporation

where

pdv <3000-32000> (continued)	On the SFP, the maximum PDV can be reached under specific circumstances. The actual PDV is about 4 ms below the maximum value. For a CSoP, the realistic achievable PDV is 12 ms, and for TPoP it is 28 ms. If the PDV is set to 10 ms, the actual realistic achievable PDV is 6 ms. The PDV is specified in terms of microseconds, from 1000 up to 32000. The TDM SFPs and module ports only accept PDV in increments of ms, so all values are rounded up to the nearest ms. For example a PDV of 11111 microseconds is converted to 12 ms. For the ports on the Combo TDM OC48 FRU module, the default value, is not user-configurable.

Copyright© 2022 Ciena® Corporation

where

tdm-profile <TDM Profile Name>	is the TDM profile name. ***Note:*** This attribute is not required virtual circuits on the ports of the Combo TDM OC48 FRU module.
cas <on | off>	indicates whether channel associated signaling is on or off. ***Note:*** This attribute is not supported for virtual circuits on ports on the Combo TDM OC48 FRU module.
payload-size <NUMBER>	iindicates the payload size in bytes. This parameter is used only for virtual circuits on the ports of the Combo TDM OC48 FRU module. It is an optional parameter. If not explicitly set the default value, and the minimum and maximum allowable values depend on the channelization, as follows: • With DS1 channelization, the minimum payload size is 96, the maximum payload size is 192, and the default payload size is 192. • With E1 channelization, the minimum payload size is 128, the maximum payload size is 256, and the default payload size is 256. • With DS3 and DS3->STS-1 channelization, the minimum payload size is 256, the maximum payload size is 1500, and the default payload size is 512. • With E3 channelization, the minimum payload size is 256, the maximum payload size is 2000, and the default payload size is 512. • With STS-1 channelization, the minimum payload size is 256, the maximum payload size is 1980, and the default payload size is 512. • With STS-3 or STM1 channelization, the minimum payload size is 783, the maximum payload size is 3000, and the default payload size is 783. • With STS-12 or STM4 channelization, the minimum payload size is 783, the maximum payload size is 3100, and the default payload size is 783. • With STS-48 or STM16 channelization, the minimum payload size is 783, the maximum payload size is 4000, and the default payload size is 783. ***Note:*** Set the payload size to 783 for an OC48 Combo TDM FRU port if that port is interoperating with a one of the following transceivers: XCVR-CEP001, XCVR-CEP002 XCVR-CEP003.

Copyright© 2022 Ciena® Corporation

Procedure 10-36
Configuring MPLS transport

To configure MPLS transport, configure:

- IP interfaces
- IGP
- tunnels
- pseudowires

For procedures, refer to *39XX/51XX Service Delivery, Aggregation and Virtualization Switches MPLS Configuration*.

Copyright© 2022 Ciena® Corporation

Procedure 10-37
Configuring Ethernet virtual circuits

Configure Ethernet virtual circuits.

Note: TDM SFPs can only use 0x8100 as the TPID value in C-TAG (single-tagged) or (S-TAG + C-TAG) (double-tagged) frames.

Step	Action
1	Create a VLAN: `vlan create vlan <vlan>`
2	Add ports to the VLAN: `vlan add vlan <vlan> port <Port list>`
3	Create an Ethernet virtual circuit: `virtual-circuit ethernet create vc <vc> vlan <vlan>`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-38
Configuring virtual switch cross-connections

Configure a virtual switch cross-connection to associate a UNI virtual circuit with an NNI virtual circuit.

The UNI virtual circuits are mapped to attachment circuits and NNI virtual circuits are mapped to Layer 2 or Layer 3 tunnels. Associating the two virtual circuits creates a point-to-point path between attachment circuits and virtual circuits into the PSN tunnel.

Note: SFP TDM attachment circuits cannot be mapped to MPLS L2 VPNs.

Step	Action
1	Create a cross-connection between a TDM virtual circuit and an Ethernet virtual circuit or MPLS virtual circuit: `virtual-switch cross-connect create xc <xc> {tdm-vc <TDM VC Name>} {eth-vc <Virtual Circuit Name>} {mpls-vc <MPLS TDM Virtual Circuit Name>}` where xc <xc> — is the name of the cross-connection. tdm-vc <TDM VC Name> — is the TDM virtual circuit name. eth-vc <Virtual Circuit Name> — is the Ethernet virtual circuit name. mpls-vc <MPLS TDM Virtual Circuit Name> — is the MPLS virtual circuit name.
2	Enable or disable the specified cross-connection: ***Note:*** The cross-connection is enabled by default. `virtual-switch cross-connect {enable|disable} xc <xconnect-name>` where xc <xconnect-name> — is the name of the cross-connection created in step 1.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-39
Configuring jitter buffer monitoring on a T1 module

Use this procedure to set jitter buffer monitoring on a T1 module.

Step	Action

To start the jitter buffer monitoring on a port and channel for a specific duration

1 Configure jitter buffer monitoring on a T1 module:

```
port tdm t1 jitter-buffer-monitoring set t1 <port-name>
channel <channel-num> duration <monitoring-duration>
```

where

port tdm <port-name>	is the TDM port name.
channel <channel-number>	is the channel set.
duration <monitoring-duration>	is the duration configured by the user. Maximum duration can be set as 1440 minutes (24 hours).

To stop the jitter buffer monitoring on a port and channel

2 Stop jitter buffer monitoring on a T1 module:

```
port tdm t1 jitter-buffer-monitoring unset t1 <port-name>
channel <channel-number>
```

where

port tdm <port-name>	is the TDM port name.
channel <channel-number>	is the channel set.

Copyright© 2022 Ciena® Corporation

To view the jitter buffer monitoring related parameters information

3 View jitter buffer monitoring related parameter information on a T1 module:

```
port tdm t1 jitter-buffer-monitoring show t1 <port-name>
channel <channel-num>
```

where

port tdm <port-name>	is the TDM port name.
channel <channel-number>	is the channel set.

To view the jitter buffer fill levels for a specified port and channel

4 View jitter buffer fill levels for a specified port and channel:

```
port tdm t1 jitter-buffer-snapshot show t1 <port-name>
channel <channel-num>
```

where

port tdm <port-name>	is the TDM port name.
channel <channel-number>	is the channel set.

Example

This example shows the monitoring related parameters and metrics calculated based on fill levels:

```
>port tdm e1 jitter-buffer-monitoring show e1 tdm channel 0
+--------------------------- JITTER BUFFER MONITORING --------------------+
| Parameter                                | Value                          |
+------------------------------------------+--------------------------------+
| Monitoring Status                        | ON                             |
| Port Name                                | tdm                            |
| Channel                                  | 0                              |
| Duration                (in minutes)     | 500                            |
| Clock Mode                               | adaptive                       |
| Number of Samples Collected              | 31/500                         |
| Monitoring Start Date/Time               | Wed Sep 23 08:47:57 2020       |
| Default JB Recenter Position             | 16                             |
| Peak Fill Levels        (in packets)     | 5(min),17(max)                 |
| Std Dev of Fill Levels  (in packets)     | 2.87(min),0.40(max)            |
| Mean PDV                (in ms)          | 1.23                           |
+------------------------------------------+--------------------------------+
```

This example shows the current jitter buffer fill levels for a specified port and channel

```
>port tdm t1 jitter-buffer-snapshot show t1 TDM4GE channel 0
+----------------- JITTER BUFFER SNAPSHOT -------------------------+
| Parameter                    | Value                              |
+------------------------------+------------------------------------+
```


Copyright© 2022 Ciena® Corporation

```
| Port Name                    | TDM4GE                                |
| Channel                      | 0                                     |
| JB Fill Level (in packets)   | 15(min),16(max)                       |
+------------------------------+---------------------------------------+
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-40
Displaying information about the FRU module in 3926

Use this procedure to display information about the FRU module in the 3926.

Step	Action
1	Display information about the module in the 3926: `module show`
2	Display information about the OC48-TDM Combo FRU module in the 3926: `module show module OC48-TDM`

—end—

Example

Here is sample of the “module show module” command.

```
>module show
+------------------- MODULE DETAILS-------------------+
| Parameter                  | Value                  |
+----------------------------+------------------------+
| Configured FRU Type        | OC48-TDM               |
| Inserted FRU Type          | OC48-TDM(power on)     |
+----------------------------+------------------------+
```

Here is sample of the module show module OC48-TDM command.

```
>module show module OC48-TDM
+------------- MODULE STATUS -------------+
| Module    | Power Status | Latch Status |
+----------+--------------+--------------+
|OC48_TDM   | On           | Closed       |
+----------+--------------+--------------+

+------------ MODULE OC48_TDM DEVICE ID ------------+
| Parameter                  | Value                 |
+----------------------------+-----------------------+
| PCA Part Num               | 170-0176-410          |
| PCA Serial Num             | B3196240              |
| Module Part Num            | 170-0176-900 /000     |
| Module Serial Num          | M9674846              |
| CLEI Code                  | CMUIAU7GAA            |
| Manufactured Date          | 11-03-2019            |
+----------------------------+-----------------------+

+------------------------ OC48_TDM MODULE STATUS ------------------------+
| Software Status   | -                                                   |
| Hardware Status   | -                                                   |
| Hardware State    | rc.combo   combo TDM FRU started                    |
| Hw Error Code     | 0                                                   |
| Removal Status    | NOT SAFE TO REMOVE                                  |
+-------------------+-----------------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 10-41
Displaying FPGA information about FRU modules

Use this procedure to display FPGA information about the Combo TDM OC48 FRU Module or the TDM FRU Module.

Step	Action

To display FPGA information for the Combo TDM OC48 FRU Module

1 Display FPGA information for the Combo TDM OC48 FRU Module:

```
software fpga show
```

2 Display software information, including FPGA information, for the Combo TDM OC48 FRU Module:

```
software show module
```

To display FPGA information for the Combo TDM FRU Module

3 Display software information, including FPGA information, for the TDM FRU Module:

```
software show module
```

—end—

Example

This example shows the output of the “software fpga show” command for a Combo TDM OC48 FRU Module.

```
> software fpga show
+------- BOARD SUPPORT FPGA --------+
| Name        | Version              |
+-------------+----------------------+
| saros       | 00.00.2D             |
+-------------+----------------------+

+---------- DATAPATH FPGA ----------+
| Name        | Version      | Module |
+-------------+--------------+--------+
| westfold    | 03.01.1C     | -      |
+-------------+--------------+--------+
+---------------- COMBO FRU FPGA ------------------+
| Name          | Version                          |
+---------------+----------------------------------+
| sarn          | 00.00.11                         |
+---------------+----------------------------------+
| norborn       | 01.02.1                          |
+---------------+----------------------------------+
| riffelhorn    | RFH_0100_05     19 Apr 20 18:10  |
+---------------+----------------------------------+
| codechip      | 2019-09-10T17:00:08+0700         |
+---------------+----------------------------------+
```


Copyright© 2022 Ciena® Corporation

This example shows the output of the “software show module” command for a Combo TDM OC48 FRU Module.

```
> software show module
+-------------------------------------------------------------------------
-----------------+
|                     COMBO FRU MODULE FPGA                      |
+----------------+-----------------------------------+----------------------
-----------------+
|FPGA Name    | Running Fpga Version         | Available Fpga Version       |
+-------------+-----------------------------+-------------------------------+
|SARN         | 00.00.11                     | 00.00.11                      |
|NORBORN      | 01.02.19                     | 01.02.19                      |
|RIFFELHORN   | RFH_0100_05  19 Apr 20 18:10 | RFH_0100_05  19 Apr 20 18:10  |
|CODECHIP     | 2019-09-10T17:00:08+0700     | 2019-09-10T17:00:08+0700      |
+-------------+-----------------------------------+-----------------------+
```

This example shows the output of the “software show module” command for a TDM Module.

```
> software show module
+------------------- TDM Module -------------------+
| Parameter                   | Value               |
+-----------------------------+---------------------+
| Running Firmware Version    | 04070005            |
| Available Firmware Version  | 04070005            |
| Firmware Operation Require  | None                |
| Firmware Update Result      | Successful          |
+-----------------------------+---------------------+
| Running Bank                | B                   |
| Running Version             | 0B00140101          |
| Running Date                | 16 Nov 18 11:38     |
+-----------------------------+---------------------+
| Available Version           | 0B00140101          |
+-----------------------------+---------------------+
| Operation Required          | None                |
| Version Compatibility       | Yes                 |
| Traffic Impacted            | No                  |
| Update Result               | Successful          |
| Update Load Progress        | 0  %                |
+-----------------------------+---------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 10-42
Displaying TDM port information

You can display

- global performance monitoring statistics for TDM ports, which are the total counts since the last system restart or last TDM mode change
- global TDM synchronization that shows protection-group information
- TDM mode, status and configuration a specified port
- details on configuration and current active alarms for a specified port

Note: There is no command to clear the values displayed by the command port tdm show port <port> statistics.

For more information about viewing and clearing TDM statistics, refer to "Performance monitoring," in *39XX/51XX Switches and Platforms Fault and Performance Management*. Ciena recommends using performance monitoring to view TDM port statistics.

Step	Action

To display global TDM configuration and status

1 Display global TDM configuration and status for all TDM ports:

```
port tdm show
```

To display the configuration and status of all TDM ports

2 Display the TDM mode, status and configuration for all TDM ports of a certain type:

```
port tdm <oc3-stm1|oc3-stm1-cem|oc12-stm4|oc12-stm4-cem|
oc48-stm16-cem|ds3-e3|e1|t1> show
```

To display the TDM mode, status and configuration of a specified OC3 port

3 Display the TDM mode, status and configuration for the specified OC3 port:

```
port tdm oc3-stm1 show oc3-stm1 <oc3-stm1>
```

where

oc3-stm1 <oc3-stm1>	is the port name.

Copyright© 2022 Ciena® Corporation

To display the TDM mode, status and configuration of a specified OC3-STM1 port on a Combo TDM OC48 FRU

4 Display the TDM mode, status and configuration for the specified OC3-STM1 port on a Combo TDM OC48 FRU module:

```
port tdm oc3-stm1-cem show oc3-stm1-cem <oc3-stm1-cem>
```

where

oc3-stm1-cem <oc3-stm1-cem>	is the port name on the Combo TDM OC48 FRU module.

To display the TDM mode, status and configuration of a specified OC12 port

5 Display the TDM mode, status and configuration for the specified port:

```
port tdm oc12-stm4 show oc12-stm4 <oc12-stm4>
```

where

oc12-stm4 <oc12-stm4>	is the port name.

To display the TDM mode, status and configuration of a specified OC12-STM4port on a Combo TDM OC48 FRU

6 Display the TDM mode, status and configuration for the specified OC12-STM4 port on a Combo TDM OC48 FRU module:

```
port tdm oc12-stm4-cem show oc12-stm4-cem <oc12-stm4-cem>
```

where

oc12-stm4-cem <oc12-stm4-cem>	is the port name on the Combo TDM OC48 FRU module.

To display the TDM mode, status and configuration of a specified OC48-STM16 port on a Combo TDM OC48 FRU

7 Display the TDM mode, status and configuration for the specified OC48-STM16 port on a Combo TDM OC48 FRU module:

```
port tdm oc48-stm16-cem show oc48-stm16-cem <oc48-stm16-
cem>
```

where

oc48-stm16-cem <oc48-stm16-cem>	is the port name on the Combo TDM OC48 FRU module.

To display the TDM mode, status and configuration of a specified DS3 port

8 Display the TDM mode, status and configuration for the specified port:

```
port tdm ds3-e3 show ds3-e3 <ds3-e3>
```

where

ds3-e3 <ds3-e3>	is the port name.

Copyright© 2022 Ciena® Corporation

To display the TDM mode, status and configuration of a specified E3 port

9 Display the TDM mode, status and configuration for the specified port:

```
port tdm ds3-e3 show ds3-e3 <ds3-e3>
```

where

ds3-e3 <ds3-e3> is the port name.

To display the TDM mode, status and configuration of a specified T1 port

10 Display the TDM mode, status and configuration for the specified port:

```
port tdm t1 show t1 <t1>
```

where

t1 <t1> is the port name.

To display the TDM mode, status and configuration of a specified E1 port

11 Display the TDM mode, status and configuration for the specified port:

```
port tdm e1 show e1 <e1>
```

where

e1 <e1> is the port name.

To display the performance monitoring statistics for a specified TDM port

12 Display performance monitoring statistics for a specified TDM port:

```
port tdm show port <port> [statistics]
```

where

port <port> is the TDM port list

statistics displays total statistics

To display global TDM synchronization

13 Display global TDM synchronization:

```
port tdm sync show
```

To display synchronization frequency protection-group information

14 Display synchronization frequency protection-group information:

```
port tdm sync frequency protection-group show
```

To display synchronization frequency protection-group information for a specific group

15 Display synchronization frequency protection-group information for a specific group:

```
port tdm sync frequency protection-group show group
<group name>
```

—end—

Copyright© 2022 Ciena® Corporation

Example

The port tdm show command provides global configuration and status as well as TDM configuration and status for all TDM ports.

Note: After a system reboot, the global state of TDM is in “Initializing” for up to 3 minutes.

This example shows the output of the port tdm show command:

```
port tdm show

+-------------------------------------------+
|    TDM Global Configuration and Status    |
+-------------------+-----------------------+
|     Parameter     |         Value         |
+-------------------+-----------------------+
| Mode              | etsi                  |
| C-TPID            | 8100                  |
+-------------------+-----------------------+

+-----------------------------------------------------+
|         TDM Module Configuration and Status         |
+-------------------+---------------------------------+
|     Parameter     |              Value              |
+-------------------+---------------------------------+
| MAC Address       | 54:c3:3e:12:a1:0b               |
| Reference1 Clock  | -                               |
| Reference2 Clock  | -                               |
| Alarm State       | None                            |
+-------------------+---------------------------------+

+-----------------------------------------------------------------------------------------------------+
|                      TDM E1 Ports Configuration and Status                                          |
+--------+-----------------+---------+-----------+-------------------------------+-----------+-----------+
| Port   |     Framing     | Line    | Loopback  |        Protection Group       | Admin     | Oper      |
|        |                 | Coding  | Status    |                               | State     | State     |
+--------+-----------------+---------+-----------+-------------------------------+-----------+-----------+
| tdm01  | unframed        | hdb3    | disabled  | not assigned                  | disabled  | down      |
| tdm02  | unframed        | hdb3    | disabled  | not assigned                  | disabled  | down      |
| tdm03  | unframed        | hdb3    | disabled  | not assigned                  | disabled  | down      |
| tdm04  | unframed        | hdb3    | disabled  | not assigned                  | disabled  | down      |
| tdm05  | unframed        | hdb3    | disabled  | not assigned                  | disabled  | down      |
| tdm06  | unframed        | hdb3    | disabled  | not assigned                  | disabled  | down      |
+--------+-----------------+---------+-----------+-------------------------------+-----------+-----------+

+-----------------------------------------------------------------------+
|            TDM E1 SFP Ports Configuration and Status                  |
+---------+---------+------------+---------+---------+---------+--------+
|  Port   |  Phys   |  Loopback  |  Clock  |  Source |  Admin  |  Oper  |
|         |  Port   |   Status   |  Mode   |   Mac   |  State  |  State |
+---------+---------+------------+---------+---------+---------+--------+
| No entries found                                                      |
+-----------------------------------------------------------------------+

+-------------------------------------------------------------------------------------------------------+
|            TDM OC3/STM1 Ports Configuration and Status                                                |
+---------+-----------+---------+--------------+------------+----------+-----------+---------+--------+
|  Port   | Channel   | Phys    | Protection   | Loopback   | Source   | Section   | Admin   | Oper   |
|         | Type      | Port    | Role         | Status     | Mac      | Trace     | State   | State  |
+---------+-----------+---------+--------------+------------+----------+-----------+---------+--------+
| No entries found                                                                                      |
+---------+-----------+---------+--------------+------------+----------+-----------+---------+--------+

+-------------------------------------------------------------------------------------+
|            Native TDM OC3/STM1 OC48-FRU Ports Configuration and Status               |
+---------+-----------+---------+--------------+------------+-----------+---------+--------+
|  Port   | Channel   | Phys    | Protection   | Loopback   | Section   | Admin   | Oper   |
|         | Type      | Port    | Role         | Status     | Trace     | State   | State  |
+---------+-----------+---------+--------------+------------+-----------+---------+--------+
| No entries found                                                                    |
+---------+-----------+---------+--------------+------------+-----------+---------+--------+
```


Copyright© 2022 Ciena® Corporation

```
+-------------------------------------------------------------------------------+
|            TDM OC12/STM4 Ports Configuration and Status                       |
+--------+-----------+--------+------------+----------+-----------+---------+--------+
|  Port  |  Channel  |  Phys  |  Loopback  |  Source  |  Section  |  Admin  |  Oper  |
|        |   Type    |  Port  |   Status   |   Mac    |  Trace    |  State  |  State |
+--------+-----------+--------+------------+----------+------- ----+---------+--------+
| No entries found                                                              |
+--------+-----------+--------+------------+----------+-----------+---------+--------+

+-------------------------------------------------------------------------------+
|            Native TDM OC12/STM4 OC48-FRU Ports Configuration and Status        |
+--------+-----------+--------+--------------+------------+-----------+---------+--------+
|  Port  |  Channel  |  Phys  |  Protection  |  Loopback  |  Section  |  Admin  |  Oper  |
|        |   Type    |  Port  |    Role      |   Status   |  Trace    |  State  |  State |
+--------+-----------+--------+--------------+------------+-----------+---------+--------+
| No entries found                                                              |
+--------+-----------+--------+--------------+------------+-----------+---------+--------+

+-------------------------------------------------------------------------------+
|            TDM OC48/STM16 Ports Configuration and Status                      |
+--------+-----------+--------+------------+----------+-----------+---------+--------+
|  Port  |  Channel  |  Phys  |  Loopback  |  Source  |  Section  |  Admin  |  Oper  |
|        |   Type    |  Port  |   Status   |   Mac    |  Trace    |  State  |  State |
+--------+-----------+--------+------------+----------+-----------+---------+--------+
| No entries found                                                              |
+--------+-----------+--------+------------+----------+-----------+---------+--------+

+-----------------------------------------------------------------------------------------+
|            Native TDM OC48/STM16 OC48-FRU Ports Configuration and Status                 |
+-----------------+----------------+------+----------+-----------------+----------+----------+
|      Port       |    Channel     | Phys | Loopback |     Section     |  Admin   |   Oper   |
|                 |     Type       | Port |  Status  |     Trace       |  State   |  State   |
+-----------------+----------------+------+----------+-----------------+----------+----------+
|      poTdm2.3_  | sts48c-vc4-16c | 2.3  | disabled | None            | enabled  | up       |
+-----------------+----------------+------+----------+-----------------+----------+----------+

+-----------------+------------+-----------+-----------+----------+----------+
|            TDM DS3/E3 Ports Configuration and Status                          |
+-----------------+------------+-----------+-----------+----------+----------+
|      Port       |  Channel   |  Framing  | Loopback  |  Admin   |   Oper   |
|                 |   Type     |           |  Status   |  State   |  State   |
+-----------------+------------+-----------+-----------+----------+----------+
|           tdm07 | ds3-e3     | unframed  | disabled  | disabled | down     |
|           tdm08 | ds3-e3     | unframed  | disabled  | disabled | down     |
|           tdm09 | ds3-e3     | unframed  | disabled  | disabled | down     |
|           tdm10 | ds3-e3     | unframed  | disabled  | disabled | down     |
+-----------------+------------+-----------+-----------+----------+----------+

+-----------------------------------------------------------------------------------+
|            TDM SFP DS3/E3 Ports Configuration and Status                          |
+-----------------+------------+-------+----------+-------------------+----------+----------+
|      Port       |  Channel   | Phys  | Loopback |      Source       |  Admin   |   Oper   |
|                 |   Type     | Port  |  Status  |       Mac         |  State   |  State   |
+-----------------+------------+-------+----------+-------------------+----------+----------+
| No entries found                                                                  |
+-----------------+------------+-------+----------+-------------------+----------+----------+
```

This example shows the output of the port tdm oc3-stm1 show command:

```
port tdm oc3-stm1 show
+-------------------------------------------------------------------------------------------------+
|                          TDM OC3 Ports Configuration and Status                                 |
+----------+------------+------+----------+-------------------+-----------------+----------+----------+
|  Port    |  Channel   | Phys | Loopback |       Port        |     Section     |  Admin   |   Oper   |
|          |   Type     | Port |  Status  |       Mac         |     Trace       |  State   |  State   |
+----------+------------+------+----------+-------------------+-----------------+----------+----------+
|  test123 | sts3c-vc4c | 2    | disabled |     not available | None            | enabled  | down     |
|      po4 | sts3c-vc4c | 4    | disabled |     not available | None            | enabled  | down     |
|    vcop1 | sts3c-vc4c | 5    | disabled |     not available | None            | enabled  | down     |
|     test | sts3c-vc4c | 7    | disabled |     not available | 1234567890      | disabled | down     |
+----------+------------+------+----------+-------------------+-----------------+----------+----------+
```


Copyright© 2022 Ciena® Corporation

This example shows the output of the port tdm oc3-stm1 show oc3-stm1 <tdm-port-name> command for a TDM port named “test”:

```
port tdm oc3-stm1 show oc3-stm1 test
+-------------------------------------------------------------+
|        Parameter         |                Value              |
+--------------------------+-----------------------------------+
| Mode                     |  oc3                              |
| Admin State              |  disabled                         |
| Operational State        |  down                             |
| Physical Port            |  7                                |
| Channelization           |  sts3c-vc4c                       |
| MAC Address              |  not available                    |
| Loopback Type            |  port                             |
| Loopback Channel         |  0                                |
| Loopback Mode            |  terminal                         |
| Loopback Enabled         |  False                            |
| Section-Trace            |  1234567890                       |
| Clock Mode               |  internal                         |
| Channels Assigned        |  0                                |
|                          |  -                                |
| Alarm State              |  None                             |
+-------------------------------------------------------------+
```

This example shows the output of the port tdm show port <tdm-port-name> statistics command for a TDM port named “test” when the global tdm mode is ANSI:

```
port tdm show port test statistics
+-------------------------------------------------------------------+
|                        ANSI TOTAL STATISTICS                      |
+-------------------------------------------+---------------------+
| Port                                      | test                |
| Operational Status                        | down                |
+-------------------------------------------+---------------------+
| Line, Code Violations                     | 0                   |
| Line, Errored Seconds                     | 0                   |
| Line, Severely Errored Seconds            | 0                   |
| Line, Unavailable Seconds                 | 0                   |
| Line, AIS                                 | 0                   |
| Section, Code Violations                  | 0                   |
| Section, Errored Seconds                  | 0                   |
| Section, Severely Errored Seconds         | 0                   |
| Section, Severely Errored Frame Seconds   | 0                   |
| Section, Loss Of Signal Seconds           | 0                   |
+-------------------------------------------+---------------------+
```


Copyright© 2022 Ciena® Corporation

This example shows the output of the port tdm show port <tdm-port-name> statistics command for a TDM port named “test123” when the global tdm mode is ETSI:

```
port tdm show port test123 statistics

+-------------------------------------------------------+
|                  ETSI TOTAL STATISTICS                |
+----------------------------------+--------------------+
| Port                             | test123            |
| Operational Status               | down               |
+----------------------------------+--------------------+
| Errored Seconds                  | 0                  |
| Severely Errored Seconds         | 0                  |
| Unavailable Seconds              | 233                |
| Available Seconds                | 0                  |
| Total Time Seconds               | 233                |
+----------------------------------+--------------------+
```

This example shows the output of the port tdm show port <tdm-port-name> command for a TDM port named “tdmt1”:

```
port tdm show port tdmt1
+-----------------------------------------------------------------+
|        Port tdmt1 Configuration and Status Information          |
+-----------------------------------------------------------------+
|        Parameter        |               Value                   |
+-------------------------+-----------------------------------+
| Mode                    |  t1                                 |
| Admin State             |  disabled                           |
| Operational State       |  down                               |
| Physical Port           |  6                                  |
| Customer TPID           |  0x8100           (Global Default)  |
| Channelization          |  ds1-e1                             |
| SFP Factory MAC Address |  not available                      |
| Port MAC Address        |  20:80:58:4a:9d:87                  |
| SFP Source MAC Address  |  not available                      |
| Loopback Type           |  ds1-e1                             |
| Loopback Channel        |  0                                  |
| Loopback Mode           |  terminal                           |
| Loopback Enabled        |  False                              |
| PRBS-Insert             |  not available                      |
| PRBS-Monitoring         |  not available                      |
| Clock Mode              |  adaptive                           |
| Clock Status            |  disabled                           |
| Framing                 |  unframed                           |
| Line Code               |  b8zs                               |
| Line Build-Out          |  133                                |
| Channels Assigned       |  0                                  |
|                         |  X                                  |
| Alarm State             |  None                               |
| Attachment circuits     |  act1                            0  |
+-------------------------+-----------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 10-43
Displaying attachment circuit information

Display attachment circuit information.

Step	Action
1	Display attachment circuit information for all or specified attachment circuits: `attachment-circuit {satop|cesop} show [ac <ac-name>]`

—end—

Example

This example shows sample output for the attachment-circuit show command.

```
attachment-circuit show
+-------------------------------------------------------------------------------------+
|                           Attachment Circuits                                       |
+---------------------+----------+--------+-------------------------+------------------+
| Name                | Type     | Port   | Channels assigned       | TDM VC           |
+---------------------+----------+--------+-------------------------+------------------+
|  AC1-VL824          | satop    | tdm05  |                         | mefAC5-VL824     |
|  AC4-VL827          | satop    | tdm08  |                         | mefAC8-VL827     |
|  AC9a-VL828         | cesop    | tdm09  | 1-23                    | mefAC9a-VL828    |
|  AC9b-VL829         | cesop    | tdm09  | 24                      | mefAC9b-VL829    |
| *cesop01ts1-4       | cesop    | tdm01  | 1-4                     |                  |
+---------------------+----------+--------+-------------------------+------------------+
NOTE: * preceding a CESoP attachment circuit name indicates it is designated
as master clock source
```


Copyright© 2022 Ciena® Corporation

Procedure 10-44
Displaying TDM port loopbacks

Display the details of TDM port loopbacks.

TDM port loopback statistics can also optionally be displayed.

Note: OCn ports on the Combo TDM OC48 FRU module do not support port loopback.

Step	Action

To display OC3 port loopback details

1 Display OC3 port loopback details:

```
port tdm oc3-stm1 loopback show oc3-stm1 <oc3-stm1>
[statistics]
```

where

oc3-stm1 <oc3-stm1>	is the OC3 port name.

To display OC3-STM1 port loopback details on a port on a Combo TDM OC48 FRU module

2 Display loopback details for an OC3-STM1 port on a Combo TDM OC48 FRU:

```
port tdm oc3-stm1-cem loopback show oc3-stm1-cem
<oc3-stm1-cem>
```

where

oc3-stm1-cem <oc3-stm1-cem>	is the port name.

To display OC12 port loopback details

3 Display OC12 port loopback details:

```
port tdm oc12-stm4 loopback show oc12-stm4 <oc12-stm4>
```

where

oc12-stm4 <oc12-stm4>	is the OC12 port name.

Copyright© 2022 Ciena® Corporation

To display OC12-STM4 port loopback details on a port on a Combo TDM OC48 FRU module

4 Display loopback details for an OC12-STM4 port on a Combo TDM OC48 FRU:

```
port tdm oc12-stm4-cem loopback show oc12-stm4-cem
<oc12-stm4-cem>
```

where

oc12-stm4-cem <oc12-stm4-cem>	is the port name.

To display OC48-STM16 port loopback details on a port on a Combo TDM OC48 FRU module

5 Display loopback details for an OC48-STM16 port on a Combo TDM OC48 FRU:

```
port tdm oc48-stm16-cem loopback show oc48-stm16-cem
<oc48-stm16-cem>
```

where

oc48-stm16-cem <oc48-stm16-cem>	is the port name.

To display DS3-E3 port loopback details

6 Display DS3-E3 port loopback details:

```
port tdm ds3-e3 loopback show ds3-e3 <ds3-e3>
```

where

ds3-e3 <ds3-e3>	is the port name.

To display E1 port loopback details

7 Display E1 port loopback details:

```
port tdm e1 loopback show e1 <e1>
```

where

e1 <e1>	is the E1 port name.

To display T1 port loopback details

8 DisplayT1 port loopback details:

```
port tdm t1 loopback show t1 <t1>
```

where

t1 <t1>	is the t1 port name.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example shows the output of the port tdm oc3-stm1 loopback show and port tdm oc3-stm1 loopback show statistics commands for an oc3-stm1 TDM port named “test”:

```
port tdm oc3-stm1 loopback show oc3-stm1 test
+-------------------------------------------------------------------------+
|                       Port test Loopback Configuration                  |
+---------------------------------+---------------------------------------+
| Parameter                       | Value                                 |
+---------------------------------+---------------------------------------+
| Loopback Mode                   | terminal                              |
| Loopback Type                   | port                                  |
| Loopback Channel                | 0                                     |
| Loopback Enabled                | False                                 |
+---------------------------------+---------------------------------------+
port tdm oc3-stm1 loopback show oc3-stm1 test statistics
+-------------------------------------------------------------------------+
|                Port test Loopback Configuration and Statistics          |
+---------------------------------+---------------------------------------+
| Parameter                       | Value                                 |
+---------------------------------+---------------------------------------+
| Loopback Mode                   | terminal                              |
| Loopback Type                   | port                                  |
| Loopback Channel                | 0                                     |
| Loopback Enabled                | False                                 |
+---------------------------------+---------------------------------------+
| Statistic                       | Value                                 |
+---------------------------------+---------------------------------------+
|   Section, Code Violations      |                                     0 |
|   Section, LOS Seconds          |                                     0 |
|   Section, Errored Seconds      |                                     0 |
|   Section, Severely Errored Secs|                                     0 |
|   Section, SEFS                 |                                     0 |
|   Line, Code Violations         |                                     0 |
|   Line, AIS Seconds             |                                     0 |
|   Line, Errored Seconds         |                                     0 |
|   Line, Severely Errored Secs   |                                     0 |
|   Line, Unavailable Seconds     |                                     0 |
|   Path, Code Violations         |                                     0 |
|   Path, Errored Seconds         |                                     0 |
|   Path, Severely Errored Secs   |                                     0 |
|   Path, Unavailable Seconds     |                                     0 |
|   Path, Failure Count           |                                     0 |
|   Path, Available Seconds       |                                     0 |
|   Path, Total Time Seconds      |                                     0 |
+---------------------------------+---------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 10-45
Displaying virtual circuit information

You can display

- all virtual circuits of type MEF8 or a specific virtual circuit of type MEF8
- all virtual circuits of type Dry Martini or a specific virtual circuit of type Dry Martini
- all virtual circuits of type MPLS or a specific virtual circuit of type MPLS

Note: The Dry Martini and MPLS encapsulation types do not apply to the Combo TDM OC48 FRU module. See the table "Combo TDM OC48 FRU module MPLS encapsulation support".

You can also display statistics for TDM virtual circuits by creating performance monitoring instances. For more information about performance monitoring instances, refer to *39XX/51XX Switches and Platforms Fault and Performance Management*.

Step	Action

To display all virtual circuits of type MEF8 or a specific virtual circuit of type MEF8

1 Display all virtual circuits of type MEF8 or a specific virtual circuit of type MEF8:

```
virtual-circuit tdm-mef8 show [vc <vc-name>]
```

To display all virtual circuits of type Dry Martini or a specific virtual circuit of type Dry Martini

2 Display all virtual circuits of type Dry Martini or a specific virtual circuit of type Dry Martini:

```
virtual-circuit tdm-dry-martini show [vc <vc-name>]
```

To display all virtual circuits of type MPLS or a specific virtual circuit of type

3 Display all virtual circuits of type MPLS or a specific virtual circuit of type MPLS:

```
virtual-circuit tdm-mpls show [vc <vc-name>]
```

—end—

Copyright© 2022 Ciena® Corporation

Example

This example shows a sample of the virtual-circuit tdm-mef8 show command:

```
virtual-circuit tdm-mef8 show
+-------------------------------------------------------------------------+
|                           Tdm MEF8 Virtual Circuits                     |
+-----------------+-------------------+-----------------+--------+---------+
| Name            | AttachmentCircuit | Cross-connect   | In Ecn | Out Ecn |
+-----------------+-------------------+-----------------+--------+---------+
| AC1-VL900       | AC1-VL900         | XC-Cir1VL900    | 11001  | 12001   |
| mefAC16-VL913   | AC16-VL913        | XC-Cir16VL913   | 11024  | 12024   |
+-----------------+-------------------+-----------------+--------+---------+
```


Copyright© 2022 Ciena® Corporation

Procedure 10-46
Displaying virtual switch information

You can display

- all virtual switch information
- all configured cross-connections or a specific cross-connection

Step	Action

To display all virtual switch information

1 Display all virtual switch information:

```
virtual-switch show
```

To display details about a specific virtual switch

2 Display details about a specific virtual switch:

```
virtual-switch show vs <vs> {statistics}
```

where

statistics displays subscriber port statistics.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 10-47
Displaying virtual switch cross-connection information

You can display configuration and status for:

- all configured cross-connections
- a specific cross-connection

Step	Action

To display all configured cross-connections

1 Display all configured cross-connections:

```
virtual-switch cross-connect show
```

To display a specific cross-connection

2 Display all configured cross-connections or a specific cross-connection:

```
virtual-switch cross-connect show xc <xc>
```

—end—

Example

This example shows a sample of the virtual-switch cross-connect show command:

```
virtual-switch cross-connect show
+-------------------------------------------------------------------------------------+
|                              Cross-Connects                                         |
+-----------------+-----------------+--------+-----------------+---------+-----------+
| Name            | UNI VC Name     | Type   | NNI VC Name     | Type    | OperState |
+-----------------+-----------------+--------+-----------------+---------+-----------+
| XC-Cir5VL824    | mefAC5-VL824    | TDM    | VCE-5-VL824     | Ethernet| up        |
| XC-Cir6VL825    | mefAC6-VL825    | TDM    | VCE-6-VL825     | Ethernet| up        |
| XC-Cir7VL826    | mefAC7-VL826    | TDM    | VCE-7-VL826     | Ethernet| up        |
| XC-Cir8VL827    | mefAC8-VL827    | TDM    | VCE-8-VL827     | Ethernet| up        |
| XC-Cir9aVL828   | mefAC9a-VL828   | TDM    | VCE-9a-VL828    | Ethernet| up        |
| XC-Cir9bVL829   | mefAC9b-VL829   | TDM    | VCE-9b-VL829    | Ethernet| up        |
+-----------------+-----------------+--------+-----------------+---------+-----------+
```


Copyright© 2022 Ciena® Corporation

Procedure 10-48
Configuring SAToP services over 802.1Q Metro Ethernet on a TDM SFP or FRU port

Configure SAToP services over 802.1q Metro Ethernet when you want to dispatch TDM bit-streams from a T1 connection over Metro Ethernet using MEF8 encapsulation or Dry-Martini encapsulation.

This figure shows a Cell Site Gateway (CSG) router connected to an Aggregator Site Gateway (ASG) router through Metro Ethernet.

Figure 10-8
CSG router connected to ASG router through Metro Ethernet

Step	Action
1	Create a SAToP attachment circuit using the specified TDM port instance: `attachment-circuit satop create ac <ac-name> port <port>`
2	Set the TDM profile: `virtual-circuit tdm create tdm-profile <TDM Profile Name> pwType satop {payload-size <NUMBER: 1-256>}`

3	Create the virtual circuit by performing one of these sub-steps.
	a. Create a Dry Martini virtual circuit:
	`virtual-circuit tdm-dry-martini create vc <vc> ac <Attachment Circuit Name> [c-vid <NUMBER: 1-4094>] [c-pcp <NUMBER: 0-7>] {peer-mac <MAC address: XX:XX:XX:XX:XX:XX>} [pdv <NUMBER: 1000-32000>] {in-label <unit 1-1048575>} {out-label <unit 1-1048575>} {tdm-profile <TDM Profile Name>}`
	b. Create a MEF8 virtual circuit:
	`virtual-circuit tdm-mef8 create vc <vc> {ac <Attachment Circuit Name>} [c-vid <NUMBER: 1-4094>] [c-pcp <NUMBER: 0-7>] {peer-mac <MAC address:`

Copyright© 2022 Ciena® Corporation

```
XX:XX:XX:XX:XX:XX>} [pdv <NUMBER: 1000-32000>] {in-ecn
<NUMBER: 1-65535>} {out-ecn <NUMBER: 1-65535>} tdm-
profile <TDM Profile Name>
```

Note: This command is blocked if the functionality is incompatible with the transceiver type. See "Transceiver-related limitations".

4 Create a VLAN:

```
vlan create vlan <vlan>
```

5 Add ports to the VLAN:

```
vlan add vlan <vlan> port <Port list>
```

6 Create an Ethernet virtual circuit:

```
virtual-circuit ethernet create vc <vc> vlan <vlan>
```

7 Create a cross-connection:

```
virtual-switch cross-connect create xc <xc> tdm-vc <TDM
VC Name> eth-vc <Virtual Circuit Name>
```

—end—

Example

This example configures SAToP over 802.1Q Metro Ethernet. The Ethernet MAC header with CVLAN 100 serves as the PSN Layer 2 tunnel through the service provider's Ethernet network. By default, the MAC address is assigned to PWE and the C-VLAN type used is 0x8100.

```
port tdm t1 set port tdm01-tdm16 framing unframed line-code ami line-build-
out 133
attachment-circuit satop create ac satopAc1 port tdm01
virtual-circuit tdm create tdm-profile satop-profile1 pwType satop payload-
size 192
virtual-circuit tdm-dry-martini create vc satop-vc1 ac satopAc1 peer-mac
00:00:01:01:01:01 in-label 4097 out-label 4097 pdv 3000 tdm-profile satop-
profile1
vlan create vlan 100
vlan add vlan 100 port 3
virtual-circuit ethernet create vc ethernet-vc1 vlan 100
virtual-switch cross-connect create xc tdmVs tdm-vc satop-vc1 eth-vc ethernet-
vc1
```


Copyright© 2022 Ciena® Corporation

Procedure 10-49
Configuring CESoP services over 802.1Q Metro Ethernet on a TDM SFP or FRU port

Configure CESoP services over 802.1Q metro Ethernet as required by the network plan.

Step	Action
1	Create a CESoP attachment circuit using the specified TDM port instance and list of channels: `attachment-circuit cesop create ac <ac-name> port <port> channel <channel-list>`
2	Create the TDM profile: `virtual-circuit tdm create tdm-profile <TDM Profile Name> pwType cesop payload-size <1-256>}`

3	Create the virtual circuit by performing one of these sub-steps.
	a. Create a Dry Martini virtual circuit:
	`virtual-circuit tdm-dry-martini create vc <vc> ac <Attachment Circuit Name> [c-vid <NUMBER: 1-4094>] [c-pcp <NUMBER: 0-7>] {peer-mac <MAC address: XX:XX:XX:XX:XX:XX>} [pdv <NUMBER: 1000-32000>] {in-label <unit 1-1048575>} {out-label <unit 1-1048575>} {tdm-profile <TDM Profile Name>}`
	b. Create a MEF8 virtual circuit:
	`virtual-circuit tdm-mef8 create vc <vc> {ac <Attachment Circuit Name>} [c-vid <NUMBER: 1-4094>] [c-pcp <NUMBER: 0-7>] {peer-mac <MAC address: XX:XX:XX:XX:XX:XX>} [pdv <NUMBER: 1000-32000>] {in-ecn <NUMBER: 1-65535>} {out-ecn <NUMBER: 1-65535>} tdm-profile <TDM Profile Name>`
	Note: This command is blocked if the functionality is incompatible with the transceiver type. See "Transceiver-related limitations".

4	Create a VLAN: `vlan create vlan <vlan>`
5	Add ports to the VLAN: `vlan add vlan <vlan> port <Port list>`
6	Create an Ethernet virtual circuit: `virtual-circuit ethernet create vc <vc> vlan <vlan>`

Copyright© 2022 Ciena® Corporation

7 Create a cross-connection:

```
virtual-switch cross-connect create xc <xc> tdm-vc <TDM
VC Name> eth-vc <Virtual Circuit Name>
```

—end—

Example

This example sets up a CESoP PW over MEF8 using 802.1Q on channel 1 and 2 of tdm01:

```
port tdm t1 set port tdm01 framing super-frame
attachment-circuit cesop create ac test port tdm01 channel 1, 2
virtual-circuit tdm create tdm-profile p1 pwType cesop payload-size 64
vlan create vlan 100
vlan add vlan 100 port 3
virtual-circuit ethernet create vc vc1 vlan 100
virtual-circuit tdm-mef8 create vc v1 ac test tdm-profile p1 peer-mac
00:00:00:00:00:05 in ecn 100 out-ecn 100
virtual switch cross-connect create xc xc1 tdm-vc v1 eth-vc vc1
```


Copyright© 2022 Ciena® Corporation

Procedure 10-50
Configuring SAToP services over Q-in-Q Metro Ethernet on a a TDM SFP or FRU port

Configure SAToP services for a Q-in-Q-based Carrier Ethernet Network as required by the network plan.

This figure shows a CSG router connected to an ASG router through Metro Ethernet - Q-in-Q.

Figure 10-9
CSG router connected to ASG router through Metro Ethernet Q-in-Q

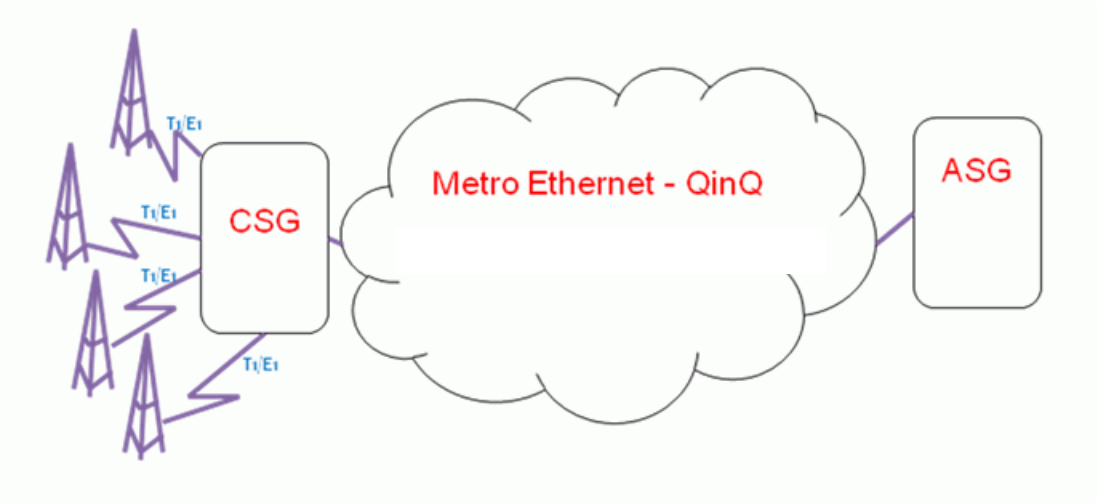

Step	Action
1	Create a SAToP attachment circuit using the specified TDM port instance: `attachment-circuit satop create ac <ac-name> port <port>`
2	Create a TDM virtual circuit by performing one of these substeps.
	a. Create a Dry Martini virtual circuit: `virtual-circuit tdm-dry-martini create vc <vc> ac <Attachment Circuit Name> [c-vid <NUMBER: 1-4094>] [c-pcp <NUMBER: 0-7>] {peer-mac <MAC address: XX:XX:XX:XX:XX:XX>} [pdv <NUMBER: 1000-32000>] {in-label <unit 1-1048575>} {out-label <unit 1-1048575>} {tdm-profile <TDM Profile Name>}`
	b. Create a MEF8 virtual circuit: `virtual-circuit tdm-mef8 create vc <vc> {ac <Attachment Circuit Name>} [c-vid <NUMBER: 1-4094>] [c-pcp <NUMBER: 0-7>] {peer-mac <MAC address:`

Copyright© 2022 Ciena® Corporation

```
XX:XX:XX:XX:XX:XX>} [pdv <NUMBER: 1000-32000>] {in-ecn
<NUMBER: 1-65535>} {out-ecn <NUMBER: 1-65535>} tdm-
profile <TDM Profile Name>
```

Note: This command is blocked if the functionality is incompatible with the transceiver type. See "Transceiver-related limitations".

3 Create a VLAN:

```
vlan create vlan <vlan>
```

4 Add ports to the VLAN:

```
vlan add vlan <vlan> port <Port list>
```

5 Create an Ethernet virtual circuit:

```
virtual-circuit ethernet create vc <vc> vlan <vlan>
```

6 Set port-based Ethernet virtual circuit attributes:

```
virtual-circuit ethernet set port <Port List> vlan-
ethertype <8100|9100|88A8> vlan-ethertype-policy
<all|encap-only|vlan-tpid>
```

7 Create a cross-connection:

```
virtual-switch cross-connect create xc <xc> tdm-vc <TDM
VC Name> eth-vc <Virtual Circuit Name>
```

—end—

Copyright© 2022 Ciena® Corporation

Example

This example configures SAToP services over Q-in-Q-based Metro Ethernet. The Ethernet MAC header with SVLAN 300 and CVLAN 100 serves as the PSN Layer 2 tunnel through the service provider's Ethernet network. By default, the MAC address is assigned to PWE and C-VLAN type used is 0x8100.

```
port tdm t1 set port tdm02 clock-mode recovered
port tdm t1 set port tdm01 framing unframed line-code ami clock-mode recovered
master-clock-src-port tdm02 line-build-out 133
attachment-circuit satop create ac satopAc1 port tdm01
virtual-circuit tdm create tdm-profile satop-profile1 pwType satop payload-
size 192
virtual-circuit tdm-mef8 create vc satopVc1 ac satopAc1 c-vid 100 peer-mac
00:00:01:01:01:01 in-ecn 100 out-ecn 200 pdv 3000 tdm-profile satop-profile1
vlan create vlan 300
vlan add vlan 300 port 3
virtual-circuit ethernet create vc ethernet-vc1 vlan 300
virtual-switch cross-connect create xc mefVs tdm-vc satopVc1 eth-vc ethernet-
vc1
```


Copyright© 2022 Ciena® Corporation

Procedure 10-51
Configuring SAToP pseudowire over an MPLS network

Configure SAToP pseudowire over an MPLS network as required by the network plan.

This figure shows a CSG router connected to an ASG router through an MPLS network.

Figure 10-10
CSG router connected to ASG router through MPLS network

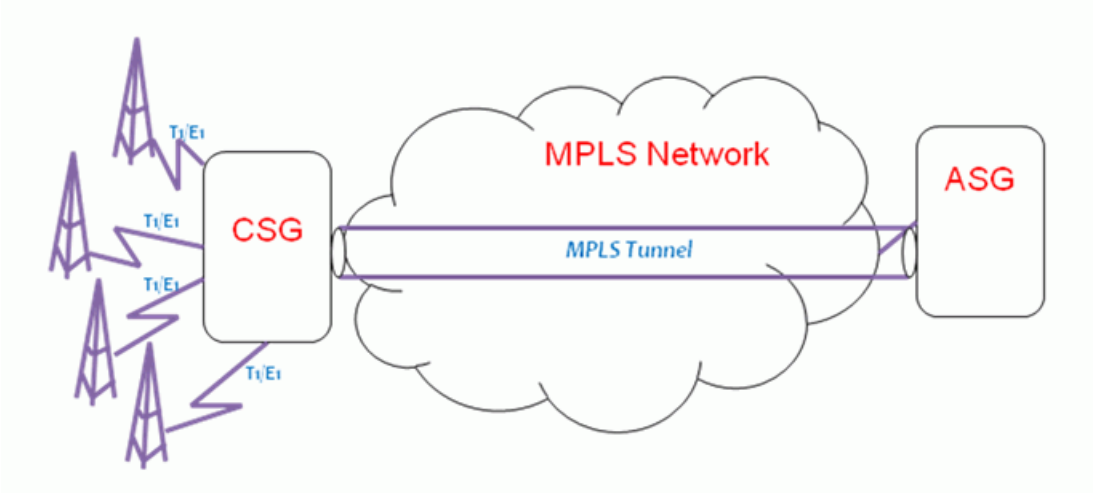

Step	Action
1	Set T1 attributes for the port: `port tdm t1 set port <Port List> framing unframed line-coding <b8zs|ami>] clock-mode recovered master-clock-src-port <TDM Port Name>] line-build-out <133|266|399|533|655>]`
2	Create a SAToP attachment circuit using the specified TDM port instance: `attachment-circuit satop create ac <ac-name> port <port>`
3	Create a TDM profile: `virtual-circuit tdm create tdm-profile <TDM Profile Name> {pwType <satop | cesop>} {payload-size <NUMBER: 1-256>}`
4	Create a TDM virtual circuit: `virtual-circuit tdm-mpls create vc <vc> ac <Attachment Circuit Name> pdv <NUMBER: 1000-32000> tdm-profile <TDM Profile Name>`

Copyright© 2022 Ciena® Corporation

5 Create static ingress TE-Tunnel:

```
mpls tunnel create static-ingress <static-ingress> dest-
ip <IP address> next-hop-ip <IP address> label <NUMBER:
16-1044479>
```

6 Create static decapsulation TE-Tunnel:

```
mpls tunnel create static-egress <static-egress> label
<NUMBER: 16-1044479> src-ip <IP address>
```

7 Create static MPLS virtual circuit:

```
mpls l2-vpn create static-vc <static-vc> pw-id <NUMBER:
1-2147483647> peer <IP address> in-label <NUMBER: 16-
1044479> out-label <NUMBER: 16-1044479> te-tunnel <MPLS
ingress primary tunnel> pwType tdm tdm-profile <TDM
Profile Name>
```

where

static-vc <static-vc>	is the static virtual circuit name.
pw-id <NUMBER: 1-2147483647>	is the VPN ID.
peer <IP address>	is the destination IP address.
in-label <NUMBER: 16-1044479>	is the MPLS decapsulation label.
out-label <NUMBER: 16-1044479>	is the MPLS encapsulation label.
te-tunnel <MPLS ingress primary tunnel>	is the ingress transport primary TE tunnel.
pwType tdm	sets the pseudowire type to TDM.
tdm-profile <TDM Profile Name>	is the TDM profile specified for the tdm virtual-circuit.

8 Create cross-connection:

```
virtual-switch cross-connect create xc <xc> tdm-vc <TDM
VC Name> mpls-vc < TDM Virtual Circuit Name>
```

—end—

Example

By default the MAC address is assigned to PWE and the C-VLAN type is 0x8100.

```
port tdm t1 set port tdm02 clock-mode recovered
port tdm t1 set port tdm01 framing unframed line-code ami clock-mode recovered
master-clock-src-port tdm02 line-build-out 133
```


Copyright© 2022 Ciena® Corporation

```
attachment-circuit satop create ac satopAc1 port tdm01
virtual-circuit tdm create tdm-profile satop-profile1 pwType satop payload-
size 192
virtual-circuit tdm-mpls create vc satop-vc1 ac satopAc1 pdv 3000 tdm-profile
satop-profile1
interface create loopback loop ip 20.20.20.20
mpls tunnel create static-ingress ingr-tnl dest-ip 2.2.2.2 next-hop-ip
10.2.2.2 out-label 500
mpls tunnel create static-egress egr-tnl in-label 600 src-ip 1.1.1.1
mpls l2-vpn create static-vc mpls-vc peer 2.2.2.2 te-tunnel ingr-tnl in-label
4097 out-label 4098 pw-id 1 pw-type tdm tdm-profile satop-profile1
virtual-switch cross-connect create xc mplsTdmVs tdm-vc satop-vc1 mpls-vc
mpls-vc
```


Copyright© 2022 Ciena® Corporation

Procedure 10-52
Configuring 1+1 linear APS groups

Use this procedure to configure 1+1 linear APS groups.

Step	Action

1	Create a 1+1 linear APS group:
	`linear-protection create protection-group <lpgroup-name> [`
	`[holdoff-time <INTEGER: 0..10000>],`
	`[revertive <on | off>],`
	`[wait-to-restore <INTEGER: 5..12>]]`
	where
	<lpgroup-name> is the 1+1 linear protection group name.
	holdoff-time <INTEGER: 0..10000> is the number of milliseconds to wait following a fault detection before switching traffic. Default is 0.
	revertive <on | off> is revertive (on) or non-revertive (off) switching. Default is off.
	wait-to-restore <INTEGER: 5..12> is the number of minutes to wait before restoring traffic to the working line after a fault has cleared. Default is 5.

2	Set one or more working parameters of a 1+1 linear APS group:
	`linear-protection set protection-group <lpgroup-name> {`
	`[holdoff-time <INTEGER: 0..10000>],`
	`[revertive <on | off>],`
	`[wait-to-restore <INTEGER: 5..12>]}`

3	Restore one or more working parameters of a 1+1 linear APS group to their default values:
	`linear-protection unset protection-group <lpgroup-name> {`
	`[holdoff-time <INTEGER: 0..10000>],`
	`[revertive <on | off>],`
	`[wait-to-restore <INTEGER: 5..12>]}`

4	Delete a 1+1 linear APS group: ***Note:*** Port members must first be removed from the protection group. `linear-protection delete protection-group <lpgroup-name>`

Copyright© 2022 Ciena® Corporation

5 Display a summary of 1+1 linear APS group configuration and status:

```
linear-protection show
```

—end—

Example

Here is sample output from the linear-protection show command:

```
> linear-protection show
+------------------------------------ LINEAR PROTECTION GROUPS -------------------------------------------+
| Protection Group    | Arch | Working Port            | Protect Port              | State    | Status      |
+---------------------+------+-------------------------+---------------------------+----------+-------------+
| protectionGroupInst | 1+1  | sampleoc12w (23; SFBER) | sampleoc12p (24; active)  | Enabled  | SF (remote) |
+---------------------+------+-------------------------+---------------------------+----------+-------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 10-53
Adding TDM ports to 1+1 linear APS groups

Use this procedure to add TDM ports to a 1+1 linear APS group.

You can also use this procedure to remove TDM ports from a linear protection group.

> ***Note 1:*** When configuring an APS 1+1 protection group on a Combo TDM OCV48 FRU module, the working and protection ports must be on the FRU.
>
> ***Note 2:*** On a Combo TDM OCV48 FRU module, APS 1+1 can be supported on the OC3/OC12 ports with CESoP as well as SAToP attachment circuits (irrespective of channelization).

Step	Action
1	Add a TDM port to a 1+1 linear APS group: `linear-protection add protection-group <lpgroup-name> {working-port | protect-port} <tdm-port-name>` where <lpgroup-name> is the 1+1 linear protection group name. {working-port | protect-port} is the role of the TDM port in the 1+1 linear protection group: • working-port—this TDM port is a working port. • protect-port—this TDM port is a protection port. <tdm-port-name> is the TDM port name
2	Remove a TDM port from a linear protection group: `linear-protection remove protection-group <lpgroup-name> {working-port | protect-port} <tdm-port-name>`
3	Display a detailed configuration and status for a specific 1+1 linear APS group: `linear-protection show protection-group <lpgroup-name>`

—end—

Copyright© 2022 Ciena® Corporation

Example

Here is sample output from the linear-protection show protection-group command:

```
> linear-protection show protection-group protectionGroupInst
+----------- LINEAR PROTECTION GROUP protectionGroupInst ------------------+
| Parameter                          | Value                                |
+------------------------------------+--------------------------------------+
| ID                                 | 1                                    |
| Working Port                       | sampleoc12w (23; SFBER)              |
| Protect Port                       | sampleoc12p (24; active)             |
| Architecture                       | 1+1                                  |
| Switching Type (Admin)             | Bidirectional                        |
| Switching Type (Oper)              | Unidirectional                       |
| Detected Mismatch                  | Architecture                         |
| Operation Mode                     | Non-Revertive                        |
| Wait To Restore Time               | 300 seconds                          |
| Wait To Restore Remaining          | 0 seconds                            |
| Holdoff Time                       | 0 milliseconds                       |
| State (Admin)                      | Enabled                              |
| Local Switch Command               | Forced, Working to Protect           |
| Last Switch Reason                 | Signal Fail (remote)                 |
| Time of Last Switchover            | 2017-01-20 16:50:10                  |
| Last K1 Byte Tx                    | No request, null channel (0x00)      |
| Last K2 Byte Tx                    | 1+1 Bidirectional, null channel (0x05) |
| Last K1 Byte Rx                    | Signal Fail, working channel 1 (0xD1) |
| Last K2 Byte Rx                    | 1:n Unidirectional, null channel (0xC0) |
+------------------------------------+--------------------------------------+
```

Here is sample output from the port tdm show command:

```
device> port tdm show
+-----------------------------------------+
|    TDM Global Configuration and Status  |
+-----------------+-----------------------+
|     Parameter   |         Value         |
+-----------------+-----------------------+
| Mode            | ansi                  |
| C-TPID          | 8100                  |
+-----------------+-----------------------+

+-----------------------------------------------------------------------------------------------+
|                                TDM T1 SFP Ports Configuration and Status                      |
+-----------------+------+--------+----------+-------------+-----------------+----------+-------+
|      Port       | Phys | Line   | Loopback |   Clock     |     Source      | Admin    | Oper  |
|                 | Port | BldOut | Status   |   Mode      |      Mac        | State    | State |
+-----------------+------+--------+----------+-------------+-----------------+----------+-------+
| No entries found                                                                              |
+-----------------+------+--------+----------+-------------+-----------------+----------+-------+

+-------------------------------------------------------------------------------------------------------+
|                                   TDM OC3/STM1 Ports Configuration and Status                         |
+-----------------+---------------+------+-----------+---------+-----------------+----------------+---------+---------+
|      Port       |   Channel     | Phys | Protection | Loopback |     Source      |    Section     | Admin   |  Oper   |
|                 |    Type       | Port | Group Role | Status   |      Mac        |    Trace       | State   |  State  |
+-----------------+---------------+------+-----------+---------+-----------------+----------------+---------+---------+
|    sampleoc3w   | ds1-e1        | 23   | working   | disabled | 9c:7a:03:7d:a1:12 | None           | enabled | up      |
|    sampleoc3p   | ds1-e1        | 24   | protect   | disabled | 9c:7a:03:7d:a1:12 | None           | enabled | up      |
+-----------------+---------------+------+-----------+---------+-----------------+----------------+---------+---------+

+-------------------------------------------------------------------------------------------------------+
|                                  TDM OC12/STM4 Ports Configuration and Status                         |
+-----------------+---------------+------+-----------+---------+-----------------+----------------+---------+---------+
|      Port       |   Channel     | Phys | Protection | Loopback |     Source      |    Section     | Admin   |  Oper   |
|                 |    Type       | Port | Group Role | Status   |      Mac        |    Trace       | State   |  State  |
+-----------------+---------------+------+-----------+---------+-----------------+----------------+---------+---------+
| No entries found                                                                                      |
+-----------------+---------------+------+-----------+---------+-----------------+----------------+---------+---------+

+-----------------------------------------------------------------------------------------+
|                          TDM DS3/E3 Ports Configuration and Status                      |
+-----------------+-------------+-----+----------+-----------------+---------+---------+
|      Port       |  Channel    | Phys | Loopback |     Source      |  Admin  |  Oper   |
```


Copyright© 2022 Ciena® Corporation

```
|                |   Type     | Port |  Status  |       Mac        |  State   |   State  |
+----------------+------------+------+----------+------------------+----------+----------+
| No entries found                                                                         |
+----------------+------------+------+----------+------------------+----------+----------+
```

Here is sample output from the port tdm show port command:

```
> port tdm show port sampleoc3p
+-------------------------------------------------------------+
|     Port sampleoc3p Configuration and Status Information    |
+-------------------------------------------------------------+
|         Parameter         |                Value            |
+---------------------------+-------------------------------+
| Mode                      | oc3                           |
| Admin State               | enabled                       |
| Operational State         | up                            |
| Physical Port             | 24                            |
| Customer TPID             | 0x8100           (Global Default)
| Channelization            | ds1-e1                        |
| SFP Factory MAC Address   | 00:19:3a:01:5c:1c             |
| Port MAC Address          | 9c:7a:03:7d:a1:13             |
| SFP Source MAC Address    | 9c:7a:03:7d:a1:12             |
| Protection Group Name     | protectionGroupInst           |
| Protection Group Role     | protect                       |
| Protection Peer Port      | sampleoc3w                    |
| Loopback Type             | port                          |
| Loopback Channel          | 0                             |
| Loopback Mode             | terminal                      |
| Loopback Enabled          | False                         |
| Section-Trace             | None                          |
| PRBS-Insert               | none                          |
| PRBS-Ins-Channel          | 0                             |
| PRBS-Monitoring           | none                          |
| PRBS-Mon-Channel          | 0                             |
| Clock Mode                | internal                      |
| Channels Assigned         |          1         2         3  |
|                           | 01234567890123456789012345678901 |
|                           | -XXXXX--------------------------- |
|                           | 3       4         5         6  |
|                           | 23456789012345678901234567890123 |
|                           | --------------------------------- |
|                           | 6     7         8             |
|                           | 45678901234567890123          |
|                           | --------------------          |
| Alarm State               | None                          |
| Attachment circuits       | acTdmP17C01                   1 |
|                           | acTdmP17C02                   2 |
|                           | acTdmP17C03                   3 |
|                           | acTdmP17C04                   4 |
|                           | acTdmP17C05                   5 |
+-------------------------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 10-54
Enabling 1+1 linear APS groups

Use this procedure to administratively enable a 1+1 linear APS group.

You can also use this procedure to administratively disable a linear protection group.

Step	Action
1	Enable a 1+1 linear APS group: `linear-protection enable protection-group <lpgroup-name>` where <lpgroup-name> is the 1+1 linear protection group name.
2	Disable a linear protection group: `linear-protection disable protection-group <lpgroup-name>`
3	Display a summary of 1+1 linear APS group configuration and status: `linear-protection show`

—end—

Example

Here is sample output from the linear-protection show command for a 1+1 linear APS group that was enabled:

```
> linear-protection show
+------------------------------------ LINEAR PROTECTION GROUPS -------------------------------------------+
| Protection Group    | Arch | Working Port            | Protect Port              | State    | Status      |
+---------------------+------+-------------------------+---------------------------+----------+-------------+
| protectionGroupInst | 1+1  | sampleoc12w (23; SFBER) | sampleoc12p (24; active)  | Enabled  | SF (remote) |
+---------------------+------+-------------------------+---------------------------+----------+-------------+
```


Copyright© 2022 Ciena® Corporation

Here is sample configuration for a 1+1 linear APS group.

> ***Note:*** For sample configuration on a Combo TDM OC48 FRU module, see the next example.

The tasks are performed in this order:

1. Configure TDM ports with their mac attributes set to “port-mac”
2. Create attachment circuit
3. Create TDM MEF8 Virtual Circuit
4. Create TP tunnel (assuming interfaces and routing are established)
5. Create L2 VPN
6. Create virtual switch
7. Attach MPLS VC to virtual switch
8. Add TDM VC to virtual switch
9. Create 1+1 linear APS group
10. Add TDM ports to 1+1 linear APS group
11. Enable 1+1 linear APS group

```
rstp disable port 13
lldp set port 13 mode disable
port tdm oc3-stm1 create oc3-stm1 tdmP13 port 13 mac port-mac
port tdm enable port tdmP13
rstp disable port 15
lldp set port 15 mode disable
port tdm oc3-stm1 create oc3-stm1 tdmP15 port 15 mac port-mac
port tdm enable port tdmP15
attachment-circuit satop create ac acTdmP13C01 port tdmP13 channel 1 channel-
type ds1-e1
virtual-circuit tdm-mef8 create vc vcTdmP13C01 ac acTdmP13C01 peer-mac
00:02:5a:01:da:71 in-ecn 101 out-ecn 101
gmpls tp-tunnel create static-ingress-corout tnlToCambie dest-ip 41.41.41.41
next-hop-ip 172.16.10.41 forward-out-label 1001 reverse-in-label 1002
mpls l2-vpn create static-vc vcPw101 pw-id 101 peer 41.41.41.41 in-label 5001
out-label 5001 tp-tunnel-ingr-corout tnlToCambie
virtual-switch ethernet attach mpls-vc vcPw101 vs vsTdm101
virtual-switch ethernet attach mpls-vc vcPw101 vs vsTdm101
virtual-switch ethernet add vs vsTdm101 tdm-vc vcTdmP13C01 vlan 101
linear-protection create protection-group apsGroup1 revertive off
linear-protection add protection-group apsGroup1 working-port tdmP13
linear-protection add protection-group apsGroup1 protect-port tdmP15
linear-protection enable protection-group apsGroup1
```


Copyright© 2022 Ciena® Corporation

Here is sample configuration for a 1+1 linear APS group on a Combo TDM OC48 FRU module.

The tasks are performed in this order:

1. Configure TDM ports.

   Note: For ports on a Combo TRM OC48 FRU module, the source MAC defaults to the Combo TDM OC48 FRU address. The “port-mac” attribute does not apply on these ports.

2. Create attachment circuit.
3. Create TDM MEF8 Virtual Circuit.
4. Create TP tunnel (assuming interfaces and routing are established).
5. Create L2 VPN.
6. Create virtual switch.
7. Create cross-connect between TDM MPLS VC and MPLS VC.
8. Create 1+1 linear APS group.
9. Add TDM ports to 1+1 linear APS group.
10. Enable 1+1 linear APS group.

```
gmpls tp-tunnel create static-ingress-corout tnlToCambie dest-ip 41.41.41.41
next-hop-ip 172.16.10.41 forward-out-label 1001 reverse-in-label 1002
mpls l2-vpn create static-vc vcPw101 pw-id 101 peer 41.41.41.41 in-label 5001
out-label 5001 tp-tunnel-ingr-corout tnlToCambie
port tdm oc12-stm4-cem create oc12-stm4-cem poTdm2.3 port 2.3
port tdm enable port poTdm2.3
port tdm oc12-stm4-cem create oc12-stm4-cem poTdm2.4 port 2.4
port tdm enable port poTdm2.4
attachment-circuit cesop create ac tdmAc port poTdm2.3 channel 0 channel-type
sts12c-vc4-4c rtp off
virtual-circuit tdm-mpls create vc tdmMplsVc ac tdmAc
virtual-switch cross-connect create xc xc1 tdm-vc tdmMplsVc mpls-vc vcPw101
linear-protection create protection-group aps revertive on
linear-protection add protection-group aps working-port poTdm2.3
linear-protection add protection-group aps protect-port poTdm2.4
linear-protection enable protection-group aps
```


Copyright© 2022 Ciena® Corporation

Procedure 10-55
Switching over 1+1 linear APS groups

Use this procedure to perform a 1+1 linear protection group switch operation.

Step	Action

1	Perform a protection switch operation:
	`linear-protection switchover protection-group <lpgroup-name> {`
	`{forced} {working-to-protect | protect-to-working} |`
	`{clear}}`

where

<lpgroup-name>	is the 1+1 linear protection group name.
{forced}	forces a switch even if error rates are above thresholds. This change is persistent (saved to the configuration).
{working-to-protect | protect-to-working}	switches traffic from the working channel to the protection channel and vice versa
{clear}	clears a previously forced switch, manual switch, or lockout.

—end—

Copyright© 2022 Ciena® Corporation

Example

Here is sample output from the linear-protection show protection-group command for a 1+1 linear protection group whose traffic was forcibly switched to working-to-protect:

```
> linear-protection show protection-group protectionGroupInst
+----------- LINEAR PROTECTION GROUP protectionGroupInst ------------------+
| Parameter                        | Value                                  |
+----------------------------------+----------------------------------------+
| ID                               | 1                                      |
| Working Port                     | sampleoc12w (23; SFBER)                |
| Protect Port                     | sampleoc12p (24; active)               |
| Architecture                     | 1+1                                    |
| Switching Type (Admin)           | Bidirectional                          |
| Switching Type (Oper)            | Unidirectional                         |
| Detected Mismatch                | Architecture                           |
| Operation Mode                   | Non-Revertive                          |
| Wait To Restore Time             | 300 seconds                            |
| Wait To Restore Remaining        | 0 seconds                              |
| Holdoff Time                     | 0 milliseconds                         |
| State (Admin)                    | Enabled                                |
| Local Switch Command             | Forced, Working to Protect             |
| Last Switch Reason               | Signal Fail (remote)                   |
| Time of Last Switchover          | 2017-01-20 16:50:10                    |
| Last K1 Byte Tx                  | No request, null channel (0x00)        |
| Last K2 Byte Tx                  | 1+1 Bidirectional, null channel (0x05) |
| Last K1 Byte Rx                  | Signal Fail, working channel 1 (0xD1)  |
| Last K2 Byte Rx                  | 1:n Unidirectional, null channel (0xC0) |
+----------------------------------+----------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 10-56
Showing and clearing 1+1 linear APS group statistics

Use this procedure to show and clear statistics for 1+1 linear APS groups.

Step	Action
1	Show statistics for a 1+1 linear APS group: `linear-protection show protection-group <lpgroup-name> statistics` where <lpgroup-name> is the 1+1 linear protection group name.
2	Clear statistics for a 1+1 linear APS group: `linear-protection clear protection-group <lpgroup-name> statistics`

—end—

Example

Here is sample output from the linear-protection show protection-group statistics command:

```
> linear-protection show protection-group protectionGroupInst statistics
+- LINEAR PROTECTION GROUP protectionGroupInst STATISTICS -+
| Parameter                          | Value                 |
+------------------------------------+-----------------------+
| Total Switchovers                  | 0                     |
| Local Switchovers                  | 0                     |
| Remote Switchovers                 | 0                     |
| APS Mismatches                     | 0                     |
| Channel Mismatches                 | 0                     |
| Far-end Protection Line Failures   | 0                     |
| Protection Switch Byte Failures    | 0                     |
| Remote Timeouts                    | 0                     |
+------------------------------------+-----------------------+
```

—end—

Copyright© 2022 Ciena® Corporation

Layer 2 Control Frame Tunneling configuration

Layer 2 Control Frame Tunneling (L2-CFT) allows the network operator to change the handling of L2 control frames so that they are processed and forwarded as if they were data frames instead of being discarded or processed locally. Additionally, there is an option to translate the destination MAC addresses of L2 control frames to a standard L2PT MAC address, which ensures that no inadvertent peering or blocking by intermediate devices can take place.

L2-CFT recognizes an untagged L2CF in its original standard protocol form as well as after is has been encapsulated for forwarding by inserting a tag (QinQ) or adding an encapsulation, for example, MPLS-TP.

Depending on the state of processing, L2 control frames can be in three forms. This table shows the frame format for the forms for each control protocol and the default disposition for each control protocol.

Table 11-1
Control frame formats and disposition

Control frame name (with CLI attribute)	Form	Default disposition	DA MAC address	Type/Length encapsulation	LLC or EtherType	Subtype
xSTP**	Untagged	Peer	01-80-C2-00-00-00	N/A	DSAP/ SSAP=0x42 Ctrl=0x03	N/A
	Transparent	Forward	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

Copyright© 2022 Ciena® Corporation

Table 11-1
Control frame formats and disposition

Control frame name (with CLI attribute)	Form	Default disposition	DA MAC address	Type/Length encapsulation	LLC or EtherType	Subtype
Link Aggregation Control Protocol (lacp)	Untagged	Peer	01-80-C2-00-00-02	N/A	88-09	01
	Transparent	Forward	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
	Transparent	Forward	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	Peer	01-80-C2-00-00-02	N/A	88-09	03
	Transparent	Forward	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
IEEE Std. 802.1X/ EAPOL(802.1x)	Untagged	Discard	01-80-C2-00-00-03	N/A	88-8E	N/A
	Transparent	Forward	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

Copyright© 2022 Ciena® Corporation

Table 11-1
Control frame formats and disposition

Control frame name (with CLI attribute)	Form	Default disposition	DA MAC address	Type/Length encapsulation	LLC or EtherType	Subtype
LLDP 802.1ab(lldp)	Untagged	Peer	01-80-C2-00-00-0E	N/A	88-CC	N/A
	Transparent	Forward	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
E-LMI	Untagged	Discard	01-80-C2--00-00-07	N/A	88-EE	N/A
	Transparent	Forward	01-80-C2--00-00-07	N/A	88-EE	N/A
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	N/A	88-EE	N/A
PTP Peer-Delay	Untagged	Discard	01-80-C2--00-00-0E	N/A	88-F7	N/A
	Transparent	Forward	01-80-C2--00-00-0E	N/A	88-F7	N/A
	L2PT	Same as data frame	01-00-0C-CD-CD-D0	N/A	88-F7	N/A
Cisco-Port Aggregation Protocol (cisco-pagp)	Untagged	Discard	01-00-0C-CD-CD-D0	SNAP	01-04	N/A
	Transparent	Forward	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

Copyright© 2022 Ciena® Corporation

Table 11-1
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-Unidirectional Link Detection (cisco-udld)	Untagged	Discard	01-00-0C-CC-CC-CC	SNAP	01-11	N/A
	Transparent	Forward	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
	Transparent	Forward	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
	Transparent	Forward	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
	Transparent	Forward	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© 2022 Ciena® Corporation

Table 11-1
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
	Transparent	Forward	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
	Transparent	Forward	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
	Transparent	Forward	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

Copyright© 2022 Ciena® Corporation

Table 11-1
Control frame formats and disposition

Control frame name (with CLI attribute)	Form	Default disposition	DA MAC address	Type/Length encapsulation	LLC or EtherType	Subtype
GARP Multicast Registration Protocol (gmrp)	Untagged	Discard	01-80-C2-00-00-20	N/A	00-01	N/A
	Transparent	Forward	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
	Transparent	Forward	01-80-C2-00-00-21	N/A	00-01	N/A
	L2PT	Not supported	Not Supported	N/A	N/A	N/A
IS-IS (isis)	Untagged	Discard	01-80-C2-00-00-14	N/A	00-01	N/A
	Transparent	Forward	01-80-C2-00-00-15	N/A	00-01	N/A
	L2PT	Not supported	Not supported	N/A	N/A	N/A
Bridge Block (bridge-block)	Untagged	Discard	01-80-C2-00-00-0X	N/A	N/A	N/A
	Transparent	Forward	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	Untagged	Discard	01-80-C2-00-00-1X	N/A	N/A	N/A
	Transparent	Forward	01-80-C2-00-00-1X	N/A	N/A	N/A
	L2PT	Not supported	Not supported	N/A	N/A	N/A

Copyright© 2022 Ciena® Corporation

Table 11-1
Control frame formats and disposition

Control frame name (with CLI attribute)	Form	Default disposition	DA MAC address	Type/Length encapsulation	LLC or EtherType	Subtype
All Bridges Block#	Untagged	Discard	01-80-C2-00-00-1x where x = {0,F} - {4,5}	N/A	N/A	N/A
	Transparent	Forward	01-80-C2-00-00-1x where x = {0,F} - {4,5}	N/A	N/A	N/A
	L2PT	Not supported	Not supported	N/A	N/A	N/A
Generic Attribute Registration Protocol (GARP) Block (garp-block)	Untagged	Discard	01-80-C2-00-00-2X	N/A	N/A	N/A
	Transparent	Forward	01-80-C2-00-00-2X	N/A	N/A	N/A
	L2PT	Not supported	Not Supported	N/A	N/A	N/A
Bridge Rsvd 0B0F*	Untagged	Discard	01-80-C2-00-00-0X where X= {B,F}	N/A	N/A	N/A
	Transparent	Forward	01-80-C2-00-00-0X where X= {B,F}	N/A	N/A	N/A
	L2PT	Not supported	01-80-C2-00-00-0X where X= {B,F}	N/A	N/A	N/A

Copyright© 2022 Ciena® Corporation

Table 11-1
Control frame formats and disposition

Control frame name (with CLI attribute)	Form	Default disposition	DA MAC address	Type/Length encapsulation	LLC or EtherType	Subtype
Bridge Rsvd 0C0D*	Untagged	Discard	01-80-C2-00-00-0X where X={C,D}	N/A	N/A	N/A
	Transparent	Forward	01-80-C2-00-00-0X where X={C,D}	N/A	N/A	N/A
	L2PT	Not supported	01-80-C2-00-00-0X where X={C,D}	N/A	N/A	N/A
ESMC	Untagged	Peer	01-80-C2-00-00-02	N/A	0x8809	0A
	Transparent	Forward	01-80-C2-00-00-02	N/A	0x8809	0A
	L2PT	Same as data frame	01-00-OCCD-CDD0	N/A	0x8809	0A
CE2 Bridge Block*	Untagged	Discard	01-80-C2-00-00-0X where X={0...F}-{B,C,D,F}	N/A	N/A	N/A
	Transparent	Forward	01-80-C2-00-00-0X where X={0...F}-{B,C,D,F}	N/A	N/A	N/A
	L2PT	Not supported	01-80-C2-00-00-0X where X={0...F}-{B,C,D,F}	N/A	N/A	N/A

Copyright© 2022 Ciena® Corporation

Note 1: * indicates that the protocols are supported in MEF CE2 mode in release 6.11 and later.

Note 2: **xSTP applies to Spanning Tree Protocols including STP, MSTP and RSTP.

Note 3: # indicates that in mef-ce2 mode, IS-IS MAC block (01:80:C2:00:00:14/01:80:C2:00:00:15) are not part of all-bridges-block, and do not inherit the disposition of all-bridges-block. Frames with an IS-IS MAC block are forwarded as a data frame in mef-ce2 mode.

Note 4: Not supported indicates that these protocols in L2PT mode are transparently tunneled if the disposition is set to forward by means of Slowpath.

L2 control frame tunneling profile

L2-CFT profiles contain protocol dispositions and priority, and the tunnel-method of forwarded untagged L2CFs. A L2-CFT profile instance is a container for the L2-CFT configuration. A maximum of 64 L2-CFT profile instances is supported. A profile instance comprises these components:

- tunnel method of forwarded untagged frames
- cos-policy and fixed priority of forwarded untagged frames
- protocol disposition list for untagged, tagged or both L2CF forms

A L2-CFT profile instance can be associated to one or more customer-facing faceplate ports. A customer-facing port can only be associated to one L2-CFT instance. A L2-CFT profile instance cannot be associated to network-facing ports.

A user-defined L2-CFT instance has an empty protocol list when it is created. You can add a protocol disposition of required protocols for either untagged, tagged or both forms.

You can specify these dispositions for untagged L2CFs:

- discard
- forward

A L2-CFT profile instance handles the L2CF frame that is associated with a protocol in its disposition list. If the protocol is not in the disposition list, then its frames are handled using the default disposition for the protocol at the ingress logical port.

Tunnel method

The tunnel method controls the classification and processing of untagged L2 control frames. You can set the tunnel method of a L2-CFT profile to one of two modes: transparent or L2PT. The default mode is transparent.

Copyright© 2022 Ciena® Corporation

L2-CFT uses a special classifier called untagged-ctrl-vs to designate the forwarding domain of a logical interface to handle the forward action on untagged L2CFs that ingress its parent logical port. To specify an untagged-ctrl-vs, the port must have a customer membership in the designated virtual switch.

This table describes when an untagged L2CP frame ingresses the device and the L2-CFT profile associated to the ingress port includes this L2 control protocol with forward untagged disposition.

Table 11-2
Untagged-ctrl-vs specification

Untagged-ctrl-vs	Description
Not specified	If untagged-ctrl-vs is absent, the frame is discarded as its forwarding-domain has not been specified.
Specified	If an untagged-ctrl-vs is designated, then the specified virtual switch forwards the untagged L2CP that ingress the logical port.

Note 1: The untagged-ctrl-vs designation applies to untagged L2CFs only when the L2-CFT operation is enabled at that ingress port, and forward untagged-disposition is configured for the specific L2 control protocol in the L2-CFT instance associated with the ingress port.

Note 2: The untagged-ctrl-vs-designation is not required for ports added as EPL members in the MPLS virtual-circuit based virtual-switch.

The untagged-ctrl-vs can be active on a port with a vs-l2-transform of i-push,e-pop (default transform). The vs-l2-transform of i-stamp;push,e-match-pop:stamp can also be used if its associated L2-CFT profile is configured in transparent untagged tunnel mode. This cannot take place of the associated L2-CFT profile is configured in L2PT untagged-tunnel mode. L2-CFT allows forward ingress disposition on untagged L2CF of any L2CP only on ports using one of the previously mentioned VS-l2-transforms. Ports configured with other VS transform modes are not able to forward untagged L2CFs. You can configure the default virtual switch L2 transform mode i-push;e-pop on a port without any restriction from its L2-CFT profile configuration.

For an enabled custom L2-CFT profile on a virtual-switch member port with a provisioned untagged disposition of an L2 control protocol in that profile, untagged L2CFs ingressing on that port and belonging to that protocol are processed according to the supplied disposition. Otherwise, untagged L2CFs are processed according to the default disposition.

Copyright© 2022 Ciena® Corporation

If L2PT tunneling mode based L2-CFT is enabled on more than one port member of the virtual switch, and all member ports have the same untagged-ctrl-vs, then L2-CFT profiles on these member ports must have the same untagged disposition for the L2 protocols.

L2PT tunneling mode is supported on QiQ based untagged-ctrl-vs. If the port's untagged-ctrl-vs is an MPLS virtual circuit-based virtual switch, L2PT mode is not supported on the L2-CFT instance attached to that port.

This figure depicts untagged control frame tunneling.

Figure 11-1
Untagged

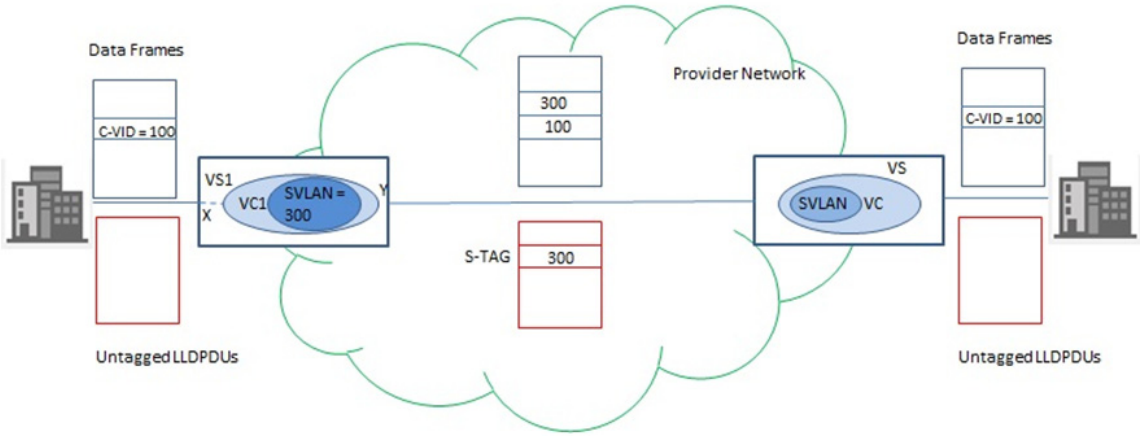

For L2 control protocols that are added to the L2-CFT instance protocol list with the untagged disposition forward, the protocols' disposition is honored for untagged L2 control frames arriving at the customer-facing interfaces. The customer-facing interfaces must be associated with this L2-CFT profile instance, have L2-CFT enabled, and have been assigned an untagged-ctrl-vs.

For L2 control protocols that are added to the L2-CFT instance protocol list with the untagged disposition discard, the protocols' disposition is honored for untagged L2 control frames arriving at the customer-facing interfaces. The customer-facing interfaces must be associated with this L2-CFT profile instance and have L2-CFT enabled.

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

Copyright© 2022 Ciena® Corporation

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.

L2-CFT profiles associated to port members of MPLS virtual switches support only the transparent mode of tunneling.

The default tunneling mode is transparent mode.

L2PT mode

When the L2-CFT profile 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/tagged-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 egress that frame out the customer-facing interfaces as a "proper" IEEE or Cisco L2 control frame, and egress the frame out the network-facing interfaces as the same L2PT frame.

> ***Note:*** If configured, L2-CFT L2PT tunneling mode works on untagged L2CFs only. Tagged L2CFs, except Cisco PVST, are always forwarded in transparent form.

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.
  - — Egressing customer-facing interface(s): The proper IEEE or Cisco MAC DA form of the frame egress out all customer-facing interfaces of the virtual switch.
  - — Egressing network-facing interface(s): The L2PT form of the frame egress out all network-facing interfaces, except for the interface on which the frame ingressed.

Copyright© 2022 Ciena® Corporation

- If the L2 control frame has an L2PT MAC DA, but is not decodable, that frame is 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 is 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 egresses 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 is 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 is 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 on the untagged-ctrl-vs attribute of that port.
  - Egressing customer-facing interface(s): The proper IEEE or Cisco MAC DA form of the frame egress out all customer-facing interfaces of the virtual switch. From customer-to-customer, the L2PT form does not propagate.
  - Egressing network-facing interface(s): The L2PT form of the frame 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 is switched out all customer-facing and network-facing interfaces with no conversion of MAC DA.

Tagged L2CF handling

In its tagged form, an L2CP frame may carry protocol-specific multicast MAC DA or L2PT MAC DA.

L2CP tagged frames handling control (discard) apply to the transparent form of tagged frames. Incoming tagged L2CP frames in L2PT form are passed without any modifications or after reverting to the original L2CP MAC DA.

Tagged L2CFs experience the same RCoS mappings and settings as those being used for the data frames associated with that service.

Copyright© 2022 Ciena® Corporation

L2-CFT's tagged L2CF handling is service unaware. You can add protocols in a L2-CFT instance with discard dispositions to L2CP tagged frames. An L2CP disposition on a tagged form is independent of an L2CP disposition on an untagged form. An L2CP disposition on untagged forms is an untagged disposition, and a tagged form is a tagged disposition. By default, tagged L2CFs are forwarded as data frames. Tagged dispositions apply to priority tagged, single tagged or double tagged frames.

> ***Note:*** Tagged L2CP controls apply to the transparent form of tagged frames only.

Double tagged L2CFs handling is supported on port members of non-QiQ virtual switches only. QiQ virtual switch port members support tagged L2CP handling on single/priority inner tagged frames only and treats other tag formats as data frames. An example of this is single frame L2CF with TPID 0x8100.

By default, all L2 control protocols have a forward disposition for tagged frames.

This figure shows an example of tagged control frame handling, where PTP Peer Delay has been given a "discard" tagged disposition in the L2-CFT profile attached to the ingress port.

Figure 11-2 Tagged

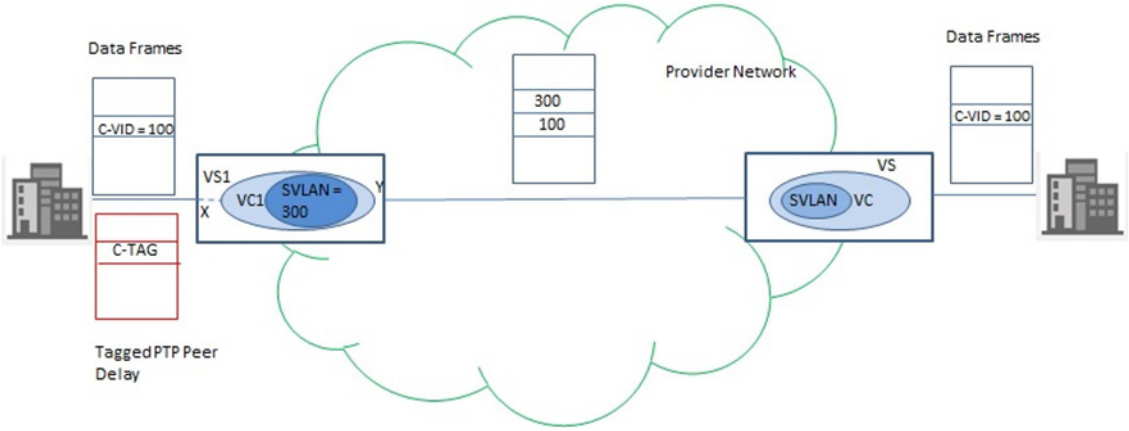

Copyright© 2022 Ciena® Corporation

Configuration

L2 control frame tunneling configuration is contained in L2-CFT profile instances. A maximum of 64 L2-CFT profile instances is supported. A control-frame tunneling profile specifies one or more protocols whose control frames are handled for untagged or tagged or both forms. For each specified protocol, the control-frame tunneling profile specifies a non-default untagged or tagged or both dispositions. If you do not specify the disposition, the default disposition of the control protocol is used. The container storing the configuration information is called an L2 control frame tunneling instance, which includes:

- Untagged CoS policy—sets the cos-policy on forwarded untagged L2 control frames. CoS policy can be fixed or ignore. Fixed untagged-cos-policy provides user-configured .1D priority to forwarded L2CFs. Ignore untagged-cos-policy does not provide differential CoS treatment to forwarded L2CFs. Instead, it allows untagged L2CFs to get CoS treatment identical to forwarded untagged data frames as they switch through the device.
- Fixed .1D priority—if “fixed” untagged-cos-policy is set on the profile, it 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 can be one of high, low and medium priorities. High, low and medium priorities can be configured at the system and they can take any value from the 0-7 range.
- A L2-CFT instance also includes tunnel-method.
- Protocol Disposition List—sets a user-defined list of protocols to be processed by L2 control frame tunneling along with untagged or tagged or both dispositions.

For the related procedure, refer to “Configuring L2 control frame tunneling” on page 11-37.

Profile processing

When a custom L2-CFT profile instance is associated to a subscriber-facing port, these attributes control L2-CFT profile processing on the port:

- L2-CFT profile processing Admin state at port level
- Port’s membership mode in untagged-ctrl-vs
- untagged-ctrl-vs type

L2-CFT profile processing must be administratively disabled at the customer-facing port before dissociating the profile from the port.

Copyright© 2022 Ciena® Corporation

This table shows how these attributes determine whether untagged dispositions and/or tagged dispositions on the L2-CFT instance include protocols are honored.

Table 11-3
Profile processing

L2-CFT Admin State at Port	Port Membership Mode in Untagged-Ctrl-VS (EVPL/EPL Don't Care)	Untagged-Ctrl-VS type (Ethernet QiQ/ MPLS/Don't Care)	Untagged-Ctrl-VS Specified (Yes/No/Don't Care)	Is L2-CFT Profile's Untagged Disposition Effective (Yes/No)		Is L2-CFT Profile's Tagged Disposition Effective (Yes/No)
				Discard	Forward	
Enabled	EPL	Ethernet QiQ	Yes	Yes	Yes	Yes
			No	Yes	No	Yes
		MPLS (VPLS/ VPWS)	Don't Care	Yes	Yes	Yes
	EVPL	Ethernet QiQ	Yes	Yes	Yes	Yes
			No	Yes	No	Yes
		MPLS (VPLS/ VPWS)	Yes	Yes	Yes	Yes
			No	Yes	No	Yes
Disabled	Don't Care	Don't Care	Don't Care	No	No	No

L2-CFT port protocol conflicts

Sometimes the configuration of port protocols such as LLDP, RSTP, OAM, LACP and LACP Marker may conflict with the configured L2-CFT protocol untagged disposition. For example, aggregated ports operate by transmitting and receiving LACP frames. It is expected that all LACP frames received on LAG ports have a forwarding disposition of PEER. A L2-CFT configuration on the aggregation port may be provisioned so that the LACP frames are dropped, resulting in a protocol disposition conflict. In most conflict cases, the L2-CFT protocol configuration takes precedence, but these port protocols are exceptions and force a forwarding disposition of PEER.

Copyright© 2022 Ciena® Corporation

This table describes the protocol conflicts for LACP and LACP Marker.

Table 11-4
Protocol conflicts

Protocol conflict	Description
LACP	Occurs when these conditions are met: • Aggregation is enabled. • L2-CFT instance attached to the LAG has added LACP to its protocol list with an untagged disposition of FORWARD or DISCARD.
LACP Marker	Occurs when these conditions are met: • Aggregation is enabled. • L2-CFT instance attached to the LAG has added LACP Marker to its protocol list with an untagged disposition of FORWARD or DISCARD. • An aggregated port was created with LACP mode of “lacp”.

A warning message is displayed when a L2-CFT disposition is overridden by a port protocol configuration.

Copyright© 2022 Ciena® Corporation

Procedures

Layer 2 Control Frame Tunneling procedures are:

- “Adding and removing untagged L2 control frame classification” on page 11-19
- “Configuring control-frame tunneling profiles” on page 11-20
- “Enabling and disabling L2 control frame tunneling” on page 11-22
- “Enabling and disabling tagged Cisco PVST L2PT forwarding” on page 11-23
- “Configuring the destination MAC address for L2PT tunnel method” on page 11-24
- “Adding and removing control protocols” on page 11-25
- “Configuring the disposition of control protocols” on page 11-27
- “Setting L2 control frame tunneling attributes” on page 11-28
- “Displaying enabled L2 control frame tunneling instances” on page 11-29
- “Displaying tagged Cisco PVST L2PT configured values” on page 11-30
- “Displaying L2-CFT resources” on page 11-31
- “Displaying the L2 control frame classification for a port” on page 11-33
- “Displaying L2 control frame tunneling configuration in the configuration file” on page 11-35
- “Configuring L2 control frame tunneling” on page 11-37
- “Configuring L2 control frame tunneling mode for MEF CE 2.0 compliance” on page 11-44

Copyright© 2022 Ciena® Corporation

Procedure 11-1
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 profile is assigned to the VS member port directly. The tunneling instance of the specified profile is used to handle the untagged L2 control frames that ingress the logical-port.

Step	Action

To add untagged L2 control frame classification on a port

1 Add tagged L2 control frame classification for a port:

```
port set port <PortName> untagged-ctrl-vs
<VirtualSwitchName>
```

To remove untagged L2 control frame classification on a port

2 Remove untagged L2 control frame classification for a port:

```
port unset port <PortName> untagged-ctrl-vs
<VirtualSwitchName>
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 11-2
Configuring control-frame tunneling profiles

You create a control-frame tunneling profile and assign it to the port on which you want to control the L2 control frames handling. The profile tells the port which types of control frames it handles.

A control-frame tunneling profile specifies one or more protocols (for example, LACP) whose control frames are handled. For each specified protocol, the control-frame tunneling profile can specify a non-default disposition. If you do not specify the disposition, the default disposition for the control protocol is used.

Possible disposition values are:

- Discard - Frame is discarded without further processing.
- Forward - frame is directed to the forwarding domain associated with the PM ingress port. In addition, these frames are flooded to all logical interfaces attached to the virtual switch associated with the ingress port.

If a port is associated to a control-frame tunneling profile, and if the port receives an untagged L2 control frame of a protocol that has not been added to that control-frame tunneling profile, that control frame is processed according to the default disposition for the control protocol.

If L2-CFT is enabled on more than one virtual switch member port, all member ports must associate with L2-CFT profile instances having the same tunneling mode.

L2-CFT profile processing must be administratively disabled at the customer-facing port before dissociating the profile from the port.

> ***Note 1:*** A maximum of 64 L2-CFT profile instances is supported.
>
> ***Note 2:*** A customer-facing port must be associated to a L2-CFT profile before enabling the L2-CFT profile processing on the port.

Copyright© 2022 Ciena® Corporation

Step	Action

To create an L2 control frame tunneling profile

1 Create an L2 control frame tunneling profile:

```
l2-cft create profile <profile> [priority
<high|low|medium>] [untagged-cos-policy <fixed|ignore>]
[tunnel-method <l2pt|transparent>]
```

where

profile <profile>	is the L2-CFT policy profile.
priority <high|low|medium>	is the egress priority for untagged frames. The default priority is high. The default values are 6 for high, 2 for low, and 4 for medium.
untagged-cos-policy <fixed | ignore>	is the policy used to assign Resolved-CoS (R-CoS) and Frame-CoS (F-CoS) priority to forwarded untagged control frames. Ignore - untagged L2CP frames are forwarded with the same priority as untagged data frames. Fixed - untagged L2CP frames are forwarded with the fixed CoS given to the profile's priority (High/Low/Medium).
tunnel-method <l2pt|transparent>	is the tunnel method for forwarded untagged L2CP frames.

To set the fixed R-CoS value for frames forwarded by an L2 control frame tunneling instance

2 Set the fixed R-CoS value for untagged L2CFs forwarded by an L2 control frame tunneling instance:

```
l2-cft set profile <profile> priority <high|low|medium>
```

Note: System level high, low and medium priorities are configurable using this command:

```
l2-cft-priority set {[high <0-7>]{low <0-7> [medium
<0-7]}
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 11-3
Enabling and disabling L2 control frame tunneling

You can enable an L2 control frame tunnel on each subscriber-facing port of each device simultaneously. The default is disabled.

The L2-CFT profile can be enabled/disabled on a port. When you disable L2 control frame tunneling profile on a port, any configured L2PT transforms do not occur, and incoming untagged and tagged L2 control frames are handled according to their default disposition.

Step	Action

To enable L2 control frame tunneling profile on a port

1 Enable L2 control frame tunneling profile on a port:

```
l2-cft enable port <port>
```

where

port <port> is the port to enable the profile.

To disable L2 control frame tunneling profile on a port

2 Disable L2 control frame tunneling profile on a port:

```
l2-cft disable port <port>
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 11-4
Enabling and disabling tagged Cisco PVST L2PT forwarding

Tagged Cisco PVST frames arriving at a customer-facing port are forwarded with L2PT MAC translation if these conditions are met:

- L2-CFT Profile has L2PT tunnel method
- L2-CFT Profile does not provide explicit tagged-disposition to Cisco PVST
- untagged-ctrl-vs is specified on the ingress port
- tagged-pvst-l2pt is enabled on the configured untagged-ctrl-vs

Step	Action

To enable tagged Cisco PVST L2PT forwarding per virtual switch

1 Enable tagged Cisco PVST L2PT forwarding per virtual switch:

```
l2-cft tagged-pvst-l2pt enable vs <vs-name>
```

where

vs <vs-name> is the virtual switch name.

To disable tagged Cisco PVST L2PT forwarding per virtual switch

2 Disable tagged Cisco PVST L2PT forwarding per virtual switch:

```
l2-cft tagged-pvst-l2pt disable vs <vs-name>
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 11-5
Configuring the destination MAC address for L2PT tunnel method

You can optionally set or remove the destination MAC address for the L2PT tunnel method.

To (optionally) set the destination MAC address for L2PT tunnel method

1 Set the destination MAC address for L2PT tunnel method:

```
l2-cft l2pt-mac set mac-name <mac-name> {mac-addr <MAC
address: XX:XX:XX:XX:XX:XX>}
```

where

mac-name <mac-name>	is the MAC name.
mac-addr <MAC address: XX:XX:XX:XX:XX:XX>}>	is the MAC address.

To (optionally) remove the destination MAC address for L2PT tunnel method

2 Remove the destination MAC address for L2PT tunnel method:

```
l2-cft l2pt-mac unset mac-name <mac-name>
```

where

mac-name <mac-name>	is the MAC name.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 11-6
Adding and removing control protocols

When adding a control protocol to a L2-CFT profile instance, you must specify untagged and/or tagged disposition of the protocol.

When you remove a control protocol from an L2 control frame tunneling instance, untagged and tagged L2 control frames are processed according to the default disposition for the control protocol.

Step	Action

To add a control protocol to an L2 control frame tunneling profile

1 Add a control protocol to an L2 control frame tunneling profile:

```
l2-cft protocol add profile <profile> {ctrl-protocol
<802.1x|all-bridges-block|bridge-block|bridge-rsvd
0B0F|bridge-rsvd-0C0D|cisco-cdp|cisco-dtp|cisco-
pagp|cisco-pvst|cisco-stp-uplink-fast|cisco-udld|cisco-
vtp|elmi|esmc|garp-block|gmrp|gvrp|lacp|lacp-
markerlldp|mef-ce2-bridge-block|oam|ptp-peer-
delay|rstp|vlan-bridge|xstp>} [untagged disposition
<discard|forward>] [tagged-disposition <discard>]
```

where

profile <profile>	is the L2-CFT profile name.
ctrl-protocol	is the control protocol.
untagged disposition <discard|forward>	Valid values are: • discard, which discards the frame without further processing • forward, which directs the frame to the forwarding domain configured as untagged-ctrl-vs of the ingress port. 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 ingress port as untagged-ctrl-vs.
tagged disposition <discard>	Valid values are: • discard, which discards the frame without further processing

Copyright© 2022 Ciena® Corporation

To remove a control protocol from an L2 control frame tunneling profile

2 Remove a control protocol from an L2 control frame tunneling profile:

```
l2-cft protocol remove profile <L2-CFT profile name>
<802.1x|all-bridges-block|bridge-block|bridge-rsvd
0B0F|bridge-rsvd-0C0D | cisco-cdp|cisco-dtp|cisco-
pagp|cisco-pvst|cisco-stp-uplink-fast|cisco-udld|cisco-
vtp|elmi|esmc|garp-block|gmrp|gvrp|lacp|lacp-
markerlldp|mef-ce2-bridge-block|oam|ptp-peer-
delay|rstp|vlan-bridge|xstp>
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 11-7
Configuring 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

To set the disposition of control protocols

1 Set the disposition of a control protocol in an L2 control frame tunneling instance:

```
l2-cft protocol set profile <profile> {ctrl-protocol
<802.1x|all-bridges-block|bridge-block|bridge-rsvd
0B0F|bridge-rsvd-0C0D | cisco-cdp|cisco-dtp|cisco-
pagp|cisco-pvst|cisco-stp-uplink-fast|cisco-udld|cisco-
vtp|elmi|esmc|garp-block|gmrp|gvrp|lacp|lacp-
markerlldp|mef-ce2-bridge-block|oam|ptp-peer-
delay|rstp|vlan-bridge|xstp>} [untagged disposition
<discard|forward>] [tagged-disposition <discard>]
```

To remove the disposition from control protocols

2 Remove the disposition from a control protocol in an L2 control frame tunneling instance:

```
l2-cft protocol unset profile <profile> {ctrl-protocol
<802.1x|all-bridges-block|bridge-block|bridge-rsvd
0B0F|bridge-rsvd-0C0D | cisco-cdp|cisco-dtp|cisco-
pagp|cisco-pvst|cisco-stp-uplink-fast|cisco-udld|cisco-
vtp|elmi|esmc|garp-block|gmrp|gvrp|lacp|lacp-
markerlldp|mef-ce2-bridge-block|oam|ptp-peer-
delay|rstp|vlan-bridge|xstp>} [untagged disposition]
[tagged-disposition]
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 11-8
Setting L2 control frame tunneling 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 high. You can modify it to a value from system-wide (high, low and medium priorities are configurable.

The tunnel method controls the classification and processing of untagged L2 control frames. You can set the tunnel method of an L2 control frame tunneling instance to one of two modes, transparent or L2PT. The default mode is transparent.

The untagged CoS policy assigns Resolved-CoS (R-CoS) and Frame-CoS (F-CoS) priority to forwarded untagged control frames. If the “ignore” policy is assigned, untagged L2CP frames are forwarded with the same priority as untagged data frames. If the “fixed” policy is assigned, untagged L2CP frames are forwarded with the same priority assigned to the profiles High, Medium and Low.

Step	Action
1	Set L2 control frame tunneling attributes: `l2-cft set profile <profile> {name <String>} [priority <high|low|medium>] [untagged-cos-policy <fixed|ignore>]{tunnel-method <l2pt|transparent>}` where profile <profile> — is the L2-CFT profile. name <String> — is the L2-CFT name. priority <high|low|medium> — is the .1D priority for encapsulated frames. The default priority is high. The default values are 6 for high, 2 for low, and 4 for medium. untagged-cos-policy <fixed | ignore> — is the policy used to assign R-CoS and F-CoS priority to forwarded untagged control frames. Ignore—untagged L2CP frames are forwarded with the same priority as untagged data frames. Fixed—untagged L2CP frames are forwarded with the same priority assigned to profiles High, Medium and Low. tunnel-method <l2pt|transparent> — is the tunnel encapsulation method. **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 11-9
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.

Step	Action
1	To display a summary of L2 control frame tunneling instances. `l2-cft show`
2	To display details of L2 control frame tunneling configuration on a port. `l2-cft show [port <port]`
3	To display details of a specific L2 control frame tunneling profile. `l2-cft show [profile <profile]`

—end—

Example

This example shows sample output for a summary of enabled L2 control frame tunneling instances.

```
l2-cft show


+------------------------- L2-CFT GLOBAL ATTRIBUTES -------------------------+
| Parameter                              | Value                              |
+----------------------------------------+------------------------------------+
| Device Mode                            | CE 1.0                             |
|-----------------------------------------------------------------------------|
+----------------------------------- PRIORITY -----------------------------------+
| High                                   | 6                                  |
| Medium                                 | 4                                  |
| Low                                    | 2                                  |
|-----------------------------------------------------------------------------|
+--------------------------- L2 CONTROL FRAME TUNNELING INFO ---------------------------+
|                                | Tunnel      | Untagged |          | Protocol | Port  |
| Profile                        | Method      | Cos      | Priority | Count    | Count |
|                                |             | Policy   |          |          |       |
+--------------------------------+-------------+----------+----------+----------+-------+
+-----------------------------------------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 11-10
Displaying tagged Cisco PVST L2PT configured values

You can display tagged Cisco PVST L2PT configured values.

Step	Action
1	Display tagged Cisco PVST L2PT configured values: `l2-cft tagged-pvst-l2pt show vs <vs_name>` where vs <vs_name> is the virtual switch name (optional).

—end—

Example

This example shows sample output for the l2-cft tagged-pvst-l2pt show command for vs1.

```
l2-cft tagged-pvst-l2pt show vs vs1

+--------- L2-CFT Tagged PVST L2PT INFO --------+
| Virtual Switch                  | Admin Status |
+---------------------------------+--------------+
| vs1                             | Enabled      |
+---------------------------------+--------------+
+--------- Virtual Switch Member Ports ---------+
| Port                            | Oper Status  |
+---------------------------------+--------------+
| 1                               | Disabled     |
+---------------------------------+--------------+
| agg1                            | Enabled      |
+---------------------------------+--------------+
```

This example shows sample output for the l2-cft tagged pvst-l2pt show command.

```
l2-cft tagged-pvst-l2pt show

+--------- L2-CFT Tagged PVST L2PT INFO --------+
| Virtual Switch                  | Admin Status |
+---------------------------------+--------------+
| vs1                             | Enabled      |
| vs2                             | Disabled     |
+---------------------------------+--------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 11-11
Displaying L2-CFT resources

You can display L2-CFT resources.

Note: The number of instances of tagged Cisco PVST L2PT forwarding support is equal to the maximum number of physical ports on the device.

Step	Action
1	Display L2-CFT resources: `l2-cft show resources` **—end—**

Example

This example shows sample output for the l2-cft show resources command.

```
l2-cft show resources

+----------- L2-CFT RESOURCE USAGE -----------+
| Parameter                   |     Count      |
|                             | Total | In Use |
+-----------------------------+-------+--------+
| L2-CFT Profiles             | 64    | 1      |
| Tagged PVST L2PT Instances  | 24    | 1      |
+-----------------------------+-------+--------+
```


Copyright© 2022 Ciena® Corporation

This example shows sample output for details about a specific enabled L2 control frame tunneling instance.

```
l2-cft show port 2

+---------------------- L2 CONTROL FRAME TUNNELING INFO -----------------------+
| Parameter             | Value                                                 |
+-----------------------+-------------------------+-------------------------+
| Port Name             | 2                                                     |
|                       |                                                       |
| Field                 | Admin                   | Oper                    |
+-----------------------+-------------------------+-------------------------+
| L2CFT Status          | Disabled                | Disabled                |
| Profile Name          |                         |                         |
| Tunnel Method         |                         |                         |
| Untagged Cos Policy   |                         |                         |
| Priority              |                         |                         |
+-----------------------+-------------------------+-------------------------+
+---------------------------------- PROTOCOLS ----------------------------------+
| Parameter                           | Value                                   |
+-------------------------------------+-----------------------------------------+
| Name                  |   Untagged Disposition    |   Tagged Disposition      |
|                       | Admin      | Oper         | Admin      | Oper         |
+-----------------------+------------+-------------+------------+-------------+
| cisco-cdp             | -          | discard     | -          | forward     |
| cisco-dtp             | -          | discard     | -          | forward     |
| cisco-pagp            | -          | discard     | -          | forward     |
| cisco-udld            | -          | discard     | -          | forward     |
| cisco-vtp             | -          | discard     | -          | forward     |
| cisco-pvst            | -          | discard     | -          | forward     |
| cisco-stp-uplink-fast | -          | discard     | -          | forward     |
| vlan-bridge           | -          | discard     | -          | forward     |
| xstp                  | -          | peer        | -          | forward     |
| lacp                  | -          | peer        | -          | forward     |
| lacp-marker           | -          | peer        | -          | forward     |
| oam                   | -          | peer        | -          | forward     |
| esmc                  | -          | peer        | -          | forward     |
| 802.1x                | -          | discard     | -          | forward     |
| elmi                  | -          | discard     | -          | forward     |
| lldp                  | -          | peer        | -          | forward     |
| ptp-peer-delay        | -          | discard     | -          | forward     |
| gmrp                  | -          | discard     | -          | forward     |
| gvrp                  | -          | discard     | -          | forward     |
| bridge-block          | -          | discard     | -          | forward     |
| all-bridges-block     | -          | discard     | -          | forward     |
| garp-block            | -          | discard     | -          | forward     |
+-----------------------+------------+-------------+------------+-------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 11-12
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. Additionally L2-CFT Status and L2-CFT Profile fields display L2-CFT status and attached profile information. **—end—**

Copyright© 2022 Ciena® Corporation

Example

This example shows sample output from the port show port command for port 1.

```
port show port 1

+---------------------------------------- PORT 1 INFO ----------------------------------------+
| Field                      | Admin                            | Oper                             |
+----------------------------+----------------------------------+----------------------------------+
| Type                       | 100/G                            | Lag (100/G)                      |
| Description                |                                  |                                  |
| Spanning Tree State        | Disabled                         |                                  |
| MAC Address                | 2c:39:c1:cc:cf:a2                |                                  |
| Phy Loopback               | Off                              |                                  |
| Link Flap Detection        | Off                              |                                  |
| Link Flap Count            | 5                                |                                  |
| Link Flap Detect Time      | 10                               |                                  |
| Link Flap Hold Time        | 300                              |                                  |
| Flow Loop Detect State     | Off                              |                                  |
| Flow Loop Detection        | Off                              | Off                              |
| L2CFT Status               | Enabled                          | Enabled                          |
| L2CFT Profile              | p111111111111111111111111111111  | t1                               |
| Link State                 | Enabled                          | Enabled                          |
| VLLI State                 | Enabled                          |                                  |
| Rx Low Power Detection     |                                  |                                  |
| State Group Link State     |                                  |                                  |
| Mode                       | default                          | sfp                              |
| Speed                      | 1 Gbps                           | 1 Gbps                           |
| Duplex                     | full                             | full                             |
| Flow Control               | off                              | off                              |
| Auto Negotiation           | Enabled                          | Enabled                          |
| Flow Control Advertised    | off                              |                                  |
| Speed Advertised           |                                  |                                  |
| Duplex Advertised          |                                  |                                  |
| PVID                       | 1                                | 1                                |
| Untag Ingress Data Vid     | 0                                | 0                                |
| Fixed Resolved CoS         | 0                                | 0                                |
| Fixed Resolved Color       | green                            | green                            |
| Ingress VLAN Filter        | Enabled                          | Enabled                          |
| Ingress VS Filter          | Off                              | Off                              |
| Acceptable Frame Type      | all                              | all                              |
| Egress Untag VLAN          | 1                                | 1                                |

| Max Frame Size             | 1526                             | 1526                             |
| Inter-Packet Gap Size      | 12                               | 12                               |
| Egress Frame Cos Policy    | ignore                           | ignore                           |
| Untagged Data VS           |                                  |                                  |
| Untagged Ctrl VS           |                                  |                                  |
| VS L2 Transform Mode       | i-push,e-pop                     |                                  |
| Resolved CoS Policy        | dot1d-tag1-cos                   | dot1d-tag1-cos                   |
| Service Port Type          | Subscriber                       | Subscriber                       |
| Eth VC EtherType           | 8100                             | 8100                             |
| Eth VC EtherType Policy    | all                              | all                              |
| Mirror-port                | Off                              | Off                              |
| Mirroring Encapsulation    | none                             |                                  |
| Mirror Encap VID           | 1                                |                                  |
| Mirror Encap TPID          | 8100                             |                                  |
| Ingress-mirror             |                                  |                                  |
| Egress-mirror              |                                  |                                  |
| Ingress to Egress QMap     | Default-RCOS                     | Default-RCOS                     |
| XCVR caps mismatch         | [NONE]                           |                                  |
| Ingress FCOS->RCOS Map     | DefaultFcosRcos                  | DefaultFcosRcos                  |
| Ingress RCOS Remark L2     | False                            | False                            |
| Egress RCOS->FCOS Map      | DefaultRcosFcos                  | DefaultRcosFcos                  |
| Benchmark State            | Disabled                         | Disabled                         |
| Ingress ACL                |                                  |                                  |
+----------------------------+----------------------------------+----------------------------------+
| Aggregation Membership     |                                  | a1                               |
| VLAN Membership            | 1000                             | 1000                             |
+----------------------------+----------------------------------+----------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 11-13
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—**

Copyright© 2022 Ciena® Corporation

Example

This example shows sample output for the configuration show command.

```
configuration show

!
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
! VIRTUAL-SWITCH MPLS AC CONFIG:
!
!
!
port set port 11 untagged-ctrl-vs vs1
!
port set port 11 untagged-ctrl-vs vs1
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
! VIRTUAL-SWITCH MPLS-VC ATTACHMENT:
!
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
! L2 CFT CONFIG:
!
l2-cft set mode mef-ce2
l2-cft l2pt-mac set mac-name default mac-addr 01:11:11:11:11:11
l2-cft create profile vs1 tunnel-method l2pt
l2-cft protocol add profile vs1 ctrl-protocol xstp untagged-disposition
discard tagged-disposition discard
l2-cft protocol add profile vs1 ctrl-protocol oam untagged-disposition discard
l2-cft protocol add profile vs1 ctrl-protocol lldp untagged-disposition
forward
l2-cft protocol add profile vs1 ctrl-protocol 802.1x untagged-disposition
discard tagged-disposition
 discard
l2-cft protocol add profile vs1 ctrl-protocol elmi untagged-disposition
discard
l2-cft protocol add profile vs1 ctrl-protocol ptp-peer-delay untagged-
disposition discard
l2-cft set port 11 profile vs1
l2-cft enable port 11
!
```


Copyright© 2022 Ciena® Corporation

Procedure 11-14
Configuring L2 control frame tunneling

In general, L2 control frame tunneling configuration includes these steps:

- Create an L2 control frame tunneling profile instance.
- Add an untagged or tagged disposition for each protocol to be handled by L2 control frame tunneling, in the profile:
  - For untagged frames, the disposition is forward or discard.
  - For tagged frames, the disposition is discard.
- Assign the L2-CFT profile to the port on which L2 frames are ingressing.
- Specify the untagged-ctrl-vs for the port on which L2 control frames are ingressing. A virtual-switch can be configured as untagged-ctrl-vs of a port only if the port is a member of that virtual-switch. This is required only if you have given forward untagged-disposition to any L2CP included in the L2-CFT profile assigned to the port.
- Enable L2-CFT instance processing on the subscriber-facing port.
- Enable L2PT tunnel-method in the profile if L2PT transforms are to occur for tunneled untagged L2 control frames. When the L2PT mode is enabled, L2CFs are processed using CPU slow-path.
- Configure the fixed .1D priority for forwarded untagged L2 control frames.
- Assign untagged-cos-policy to profile.

Note: For L2 control frame tunneling over MPLS EPL services, you do not need to specify the untagged-ctrl-vs for each virtual-switch EPL member

Copyright© 2022 Ciena® Corporation

Step	Action
1	Create the VLAN. `vlan create vlan <VlanIdList> [name <String[31]>]` where vlan <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 is “VLAN1000”. ***Note:*** VLAN names cannot begin with an integer. For example, 3PUD is not valid, but PUD3 is allowed.
2	Add network-facing ports to the VLAN. `vlan add vlan <VlanIdList> port <PortNameList>` where vlan <VlanIdList> — is the VLAN created in step 1. port <PortNameList> — is the list of ingress and egress ports.
3	To confirm the creation of the VLAN, enter this command: `vlan show`
4	Confirm that the ports are members of the VLAN. `vlan show vlan <VlanIdList>` where vlan <VlanIdList> — is the VLAN created in step 1.
5	Create a virtual circuit and assign it to the VLAN. `virtual-circuit ethernet create vc <vc> vlan <VlanIdList>` where vc <vc> — is the name of the virtual channel. vlan <VlanIdList> — is the VLAN created in step 1.
6	To confirm the creation of the virtual circuit, enter this command: `virtual-circuit show`
7	To confirm that the VLAN is associated with the virtual circuit, enter this command: `virtual-circuit show vc <VirtualCircuitName>`

Copyright© 2022 Ciena® Corporation

8 Create the virtual switch.

```
virtual-switch ethernet create vs <vs> vc <vc>
```

where

vs <vs>	is the name of the virtual switch.
vc <vc>	is the name of the virtual circuit to be associated with the virtual switch.

9 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>]
```

where

vs <vs>	is the virtual switch created in step 8.
port <PortNameList>	is the port name

10 Create L2 control frame tunneling profile and include the protocols which you want to handle differently from the default.

```
l2-cft create profile <profile> [priority
<high|low|medium>][untagged-cos-policy <fixed|ignore>]
[tunnel-method <l2pt|transparent>]
```

where

profile <profile>	is the L2-CFT policy profile.
priority <high|low|medium>]	is the egress priority for untagged L2CP frames. The default priority is high. The default values are 6 for high, 2 for low, and 4 for medium.
untagged-cos-policy <fixed | ignore>	is the policy used to assign Resolved-CoS (R-CoS) and Frame-CoS (F-CoS) priority to forwarded untagged control frames. Ignore - untagged L2CP frames are forwarded with the same priority as untagged data frames. Fixed - untagged L2CP frames are forwarded with the fixed CoS given to the profile's priority (High/Low/Medium).
[tunnel-method <l2pt|transparent>]	is the tunnel method for untagged L2CP forwarded frames.

Copyright© 2022 Ciena® Corporation

11 Assign the L2-CFT profile to the virtual switch member port on which the control frames are ingressing.

```
l2-cft set port <port> {profile <profile>}
```

where

lport <port>	is the port to set the profile on.
profile <profile>	is the L2-CFT profile.

12 Add the control protocols to the protocol disposition list with the desired disposition for the L2 control frame tunneling instance.

```
l2-cft protocol add profile <profile> {ctrl-protocol
<list>} [untagged disposition <discard|forward>] [tagged-
disposition <discard>]
```

where

port <port>	is the L2-CFT port.
{ctrl-protocol <list>}	is the control protocol to add.
[untagged disposition <discard|forward>] [tagged-disposition <discard>]	is the disposition. If left unspecified, the disposition is set to the default disposition for the specified control protocol.

13 Verify the L2-CFT profile configuration.

```
l2-cft show profile <profile>
```

14 Enable L2 control frame tunneling for the port.

```
l2-cft enable port <port>
```

where

port <port>	is the L2-CFT port.

15 Add the L2 control frame classification to the port for ingress traffic.

```
port set port <PortNameList> untagged-ctrl-vs <vs>
```

where

port <PortNameList>	is the port created in step 9.
untagged-ctrl-vs <vs>	is the virtual switch created in step 8.

16 To confirm the L2 control frame tunnel status on the given port:

```
l2-cft show port <port>
```

where

port <port>	is the L2-CFT port.

—end—

Copyright© 2022 Ciena® Corporation

Example

This 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 6
```

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© 2022 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             |                              |
+-------------------------------------------------------+
```

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
```

Add the L2 control frame classification to the port for ingress traffic.

```
port set port 1 {acceptable-frame-type all} untagged-ctrl-vs
cft-vs
```

Create the L2 control frame tunneling profile, and include the protocols which you want to handle differently from the default.

```
l2-cft create profile profile1 priority low tunnel-method l2pt
```

Assign the L2-CFT profile to the virtual switch member on which the control frames are ingressing.

```
l2-cft set port 1 profile profile1
```

Add the control protocols to the protocol disposition list with the desired disposition for the L2 control frame tunneling instance.

```
l2-cft protocol add profile profile1 elmi untagged disposition
forward
```

Add one more protocol with tagged-disposition:

```
l2-cft protocol add profile profile1 ctrl-protocol lldp
untagged-disposition discard tagged-disposition discard
```


Copyright© 2022 Ciena® Corporation

Enable L2 control frame tunneling for the port.

```
l2-cft enable port 1
```

Confirm the L2 control frame tunnel configuration for a L2-CFT profile:

```
l2-cft show profile profile1
```

```
+---------------------- L2 CONTROL FRAME TUNNELING INFO ------------------------+
| Parameter                            | Value                                   |
+--------------------------------------+-----------------------------------------+
| Profile Name                         | profile1                                |
|                                      |                                         |
| Tunnel Method                        | l2pt                                    |
|                                      |                                         |
| Untagged Cos Policy                  | fixed                                   |
|                                      |                                         |
| Priority                             | low                                     |
+--------------------------------------+-----------------------------------------+
+---------------------------------- PROTOCOLS -----------------------------------+
| Name                                 |              Disposition                |
|                                      | Untagged          | Tagged              |
+--------------------------------------+-------------------+---------------------+
| lldp                                 | discard           | discard             |
| elmi                                 | forward           | -                   |
+--------------------------------------+-------------------+---------------------+
+------------------------------- PORT ATTACHMENTS -------------------------------+
| Name                                 | Status                                  |
+--------------------------------------+-----------------------------------------+
| 1                                    | Enabled                                 |
+--------------------------------------+-----------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 11-15
Configuring L2 control frame tunneling mode for MEF CE 2.0 compliance

Configure L2 control frame tunneling mode for MEF Carrier Ethernet (CE) 2.0 compliance.

This table describes the behavior changes that occur when a device configured for CE 2.0 mode is configured for CE 1.0 mode.

Table 11-5
CE 2.0 mode to CE 1.0 mode behavior changes

Case	CE 2.0 Mode			CE 1.0 mode
	CE2-Bridge-block	Bridge-rsvd-0B0F, Bridge-rsvd-0C0D		Bridge-Block
1	forward	forward	—>	forward
2	discard	discard	—>	discard
3	forward	discard	—>	forward
4	discard	forward	—>	discard

Note: If L2 control frame tunneling is set to CE 2.0 mode and the device is downgraded to a release that does not support CE 2.0 mode, L2 control frame tunneling reverts to a CE 1.0 configuration with the behavior changes noted in the CE 2.0 mode to CE 1.0 mode behavior changes table.

Step	Action
1	Set the L2 control frame tunneling mode: `l2-cft set mode <mef-ce1|mef-ce2>` where mode <mef-ce1|mef-ce2> is the mode. **—end—**

Copyright© 2022 Ciena® Corporation

Broadcast containment

Broadcast containment prevents services from being disrupted by broadcast storms on specified ports as well as Denial of Service (DoS) attack. A broadcast storm occurs when broadcast packets flood the network, creating excessive traffic and degrading performance. When a broadcast filter is created and one or more ports are added to it, the device monitors incoming broadcast packets for a specified period of time. Broadcast packets exceeding the configured threshold are dropped. SNMP traps are sent when the configured broadcast threshold is exceeded and when broadcast traffic goes below the configured threshold.

The system software supports containment rates for these flood traffic types:

- broadcast
- unknown unicast
- unknown L2 multicast
- unknown IP multicast
- all types of multicast

Note 1: The bandwidth calculation of the broadcast containment filter may or may not include 20 bytes of the Inter Frame Gap (IFG) depending on the bandwidth calculation setting. For more information, refer to "Hardware resource management" on page 4-1.

Note 2: IPv4 multicast frames (whose destination MAC address is in the range of 01-00-5E-00-00-00 to 01-00-5E-7F-FF-FF), and IPv6 multicast frames (whose destination MAC address is in the range of 33-33-00-00-00-00 to 33-33-FF-FF-FF-FF) are not considered Layer 2 multicast frames, and are not included in the l2-mcast classification of a broadcast containment filter.

The system software has a resource mode. When resource mode is enabled, broadcast containment requires resources to be allocated with the resource manager.

Copyright© 2022 Ciena® Corporation

The default broadcast containment mode is for resource mode to be enabled. This mode is backwards compatible with previous releases (before SAOS 6.10) that did not support the non-resource mode.

When the resource mode is enabled, these rules are in place:

- All existing traffic types are supported.
- A broadcast containment resource slice is required.
- No traffic profiles can be used as traffic profiles can cause some or all traffic ingressing the port to be dropped. If traffic profiles are required, disable resource mode.
- A filter can support multiple traffic types.
- Only one filter can be applied to a port.
- Physical and aggregated ports are supported.
- The classifier combination of a single port and a virtual switch (VS) is supported. VPLS VS is supported for now.
- Dropped packets and bytes statistics are supported.
- When the filter bandwidth is set to zero, no frames are transmitted.
- Broadcast containment filters can be applied on
  - Physical or aggregation ports
  - Per port per VPLS VS
- Broadcast containment filters provide the ability to configure multiple containment classification in a filter. This table provides details:

Table 12-1
Supported broadcast containment classification types

Member type	Broadcast containment classification type			
	Unknown Unicast	**Broadcast**	**L2 Multicast**	**IP Multicast**
Port(s)	Supported	Supported	Supported	Supported
Port + VPLS VS	Supported	Supported	Supported	Not supported*
* In the case of VPLS VS, IP Multicast is accounted for as L2 Multicast.				

When the resource mode is disabled, these rules are in place:

- Three traffic types are supported: broadcast, unknown unicast, and multicast.
- A broadcast containment resource slice does not need to be allocated.
- The filter does not consume tokens from traffic profiling.
- A filter can only support one traffic type.

Copyright© 2022 Ciena® Corporation

- Multiple filters can be applied to a port.
- Only physical ports are supported. Filters must be applied individually to aggregated ports.
- Statistics are not supported.
- When the filter bandwidth is set to zero, a minimal rate of frames are transmitted, typically 64 kbps.
- Broadcast containment filters can be applied only on a physical port.
- Broadcast containment filters provide the ability to configure only a single containment classification in a filter.

Note: To enable broadcast containment commands, the Advanced Ethernet feature license must be installed with the `software license install` command.

Feature interactions

The following table lists the interactions of the broadcast containment feature.

Table 12-2
Broadcast containment feature interactions

Feature	Behavior
ACL	Frames dropped by ACL “deny” rules consume broadcast containment meter credits when broadcast containment has been configured with resource mode on. Frames discarded by broadcast containment filters are discarded regardless of any configured ACL “allow” rules. Broadcast containment statistics increment even if the frame was dropped by ACL, and vice versa.
Packet capture	Broadcast containment filters have no impact on frames “captured” by Packet Capture ingress profiles.
IP-ACL	Broadcast containment filters have no impact on frames “allowed” or “dropped” by IP-ACL rules.
Traffic profiling	If traffic profilling and broadcast containment both qualify on same type of traffic, then both consume credits of their respective meters and the rate is undeterministic.
L2 control frame tunneling	A frame forwarded in “transparent” mode using L2 control frame tunneling will be rate limited by broadcast-containment. Broadcast-containment statistics increment even if the frame was dropped or forwarded in l2pt mode by the L2 frame tunneling “discard” disposition or “l2pt” tunnel-method.
Port mirroring	Broadcast containment filters have no impact on mirrored traffic.

Copyright© 2022 Ciena® Corporation

Table 12-2
Broadcast containment feature interactions

Feature	Behavior
Benchmark	Broadcast containment filters will not be assigned to the traffic locally generated or reflected by Benchmark.
RAPS	RAPS frame destined for any other device will be rate limited by broadcast containment filters. SAOS provides a functionality to “ignore-raps-msgs” to avoid RAPS frame getting rate limited by broadcast containment filters.
TWAMP	Broadcast containment filters have no impact on TWAMP traffic destined to the device.
IGMP	Broadcast containment filters have no impact on IGMP frames destined to the device.
CFM	For 3942/5142-931/5160 – Non-Accelerated Continuity Check Messages (CCM) will be dropped by broadcast containment filters. For 392x – Non-Accelerated CCM will be dropped by broadcast containment filters if system port is reserved for the “Enhanced-CFM” feature. For 390x/393x/5142-930 –Non-Accelerated CCM will not be dropped by broadcast containment filters. On devices that support Accelerated CCM, CCM will not be dropped by broadcast containment filters.

Scale management

Scale management is as follows.

Non-resource mode

In non-resource mode, the scale of filters and its members is equal to (number of ports on device) * 3.

Resource mode

In resource mode, the scale of filters and its members are determined based on the classifiers allocated by the user through the resource-manager for the broadcast-containment feature.

The following table specifies the platform-specific scale of broadcast filters and its members:

Copyright© 2022 Ciena® Corporation

Table 12-3
Platform-specific scale of broadcast filters and members
Table 12-4

Platform	Classifiers allocated (N)	Filters that can be configured	Members that can be configured
3942, 5142, 5160	Classifiers allocated through the resource manager	N (See Note.)	One classifier required per member per containment-classification (broadcast, unknown unicast, and all multicast) added to the filter.
3924, 3926, 3928		256	
Note: N stands for the number of classifiers allocated for the broadcast-containment feature via the resource-manager.			

Note 1: For 3903, 3903x, 3904, 3905, and 3906 devices in resource mode, the scale of filters and its members is equal to (number of ports on device) * 3.

Note 2: One classifier resource is reserved for all platforms to avoid ERPS frames becoming rate-limited.

To support broadcast containment over port and LAG in resource-mode, a certain number of classifiers are reserved by default for the platforms. The number of reserved classifiers = (number of ports + number of LAGs) * 4 (types of containment classification)

This table specifies the number of classifiers reserved by default for each platform:

Table 12-5
Classifiers reserved by default for each platform

Platform	If classifiers allocated are greater than	Reserved classifiers
3924	256	76
3926	256	76
3928	256	96

Copyright© 2022 Ciena® Corporation

Limitations

Limitations to the configuration of broadcast containment per port per VPLS VS include:

- The configuration of “unknown-ip-mcast” containment classification is not supported on a filter that has a member with a classifier combination of port and VPLS VS.
- The addition of the classifier combination of port and VPLS VS is not supported on a filter with the “unknown-ip-mcast” containment classification.
- A broadcast containment filter with the classifier combination of port and VPLS VS does not allow the addition of only port as a classifier to the filter and vice versa.

Broadcast containment procedures and configuration examples are:

- “Creating a broadcast containment filter” on page 12-7
- “Setting broadcast containment filter attributes” on page 12-9
- “Setting broadcast containment configuration” on page 12-11
- “Enabling and disabling broadcast containment” on page 12-12
- “Deleting a broadcast containment filter” on page 12-13
- “Adding ports to a broadcast containment filter” on page 12-14
- “Removing a port from a broadcast containment filter” on page 12-15
- “Displaying broadcast containment filters” on page 12-16
- “Displaying a broadcast containment filter to which per port per VS membership is attached” on page 12-22
- “Displaying broadcast containment configuration in the configuration file” on page 12-23
- “Displaying all broadcast containment resources” on page 12-24
- “Configuring broadcast containment filters on a port and classifier combination of port and virtual switch” on page 12-25
- “Clearing broadcast containment statistics” on page 12-27

Copyright© 2022 Ciena® Corporation

Procedure 12-1
Creating a broadcast containment filter

The number of filters supported for each platform depends on the hardware resources allocated for the broadcast containment feature. For additional information about managing hardware resources, refer to "Hardware resource management" on page 4-1.

Step	Action
1	Create a broadcast containment filter: `broadcast-containment create filter <filter> [kbps <NUMBER: 0-10000000>] [port <Port list>] {containment-classification <all-mcast | bcast | unknown-l2-mcast | unknown-ip-mcast | unknown-ucast>} {vs <name> port <name>}` ***Note:*** The parameters `port <Port list>` and `vs <name> port <name>` are mutually exclusive.

where

filter <filter>	is the name of the broadcast containment filter.
kbps <NUMBER: 0-10000000>	is the rate in kilobits. The configured rate normalizes to the rate equal to or greater than the configured value in 64 kilobit increments to be processed by the hardware.

Copyright© 2022 Ciena® Corporation

where

containment-classification <all-mcast | bcast |unknown-l2-mcast | unknown-ip-mcast | unknown-ucast>	is the list of traffic classes covered by the filter. Valid values are: • all-mcast: filters all types of multicast packets, including known. Not supported when resource-mode is on. • bcast: filters all broadcast packets • unknown-l2-mcast: filters all unknown L2 multicast packets — L2-only devices: includes all unknown multicast frames, including unknown IP multicast frames — L3-capable devices: includes all unknown multicast frames except those with IP multicast addresses • unknown-ip-mcast: filters all unknown IP multicast packets. Not supported when resource-mode is off. — In case of MPLS configurations on a port or a MPLS virtual-switch, IP multicast frames are treated as L2 multicast frames. — Cannot be configured on a filter which has VPLS VS added as classifier. • unknown-ucast: filters all unknown unicast packets By default, the filtered traffic types are Layer 2 Broadcast traffic (bcast) and Layer 2 Unknown Multicast traffic (l2-mcast). ***Note:*** If you downgrade to a software release prior to 6.5, any containment classification settings you configure are lost, and filtered traffic types revert to Layer 2 Broadcast traffic (bcast) and Layer 2 Unknown Multicast traffic (l2-mcast).
vs <name>	specifies the virtual switch for per port per VS membership
port <name>	specifies the port for per port per VS membership

—end—

Example

This example filters all unknown multicast traffic. Note that this example assumes that the resource-mode option is enabled.

```
broadcast-containment create filter MyFilter containment-
classification unknown-l2-mcast,unknown-ip-mcast
```


Copyright© 2022 Ciena® Corporation

Procedure 12-2
Setting broadcast containment filter attributes

Set broadcast containment filter attributes to update broadcast containment filter attributes.

Step	Action
1	Set broadcast containment filter attributes: `broadcast-containment set filter <filter> [kbps <NUMBER: 0-10000000>] [name <String>] {containment-classification <all-mcast | bcast |unknown-l2-mcast | unknown-ip-mcast | unknown-ucast>}` where filter <filter> — is the name of the broadcast containment filter. kbps <NUMBER: 0-10000000> — is the rate in kilobits. The configured rate normalizes to the rate equal to or greater than the configured value in 64 kilobit increments to be processed by the hardware. port <Port list> — is the list of ports to apply the filter to.

Copyright© 2022 Ciena® Corporation

where

containment-classification <all-mcast | bcast |unknown-l2-mcast | unknown-ip-mcast | unknown-ucast>	is the list of traffic classes covered by the filter. Valid values are: • all-mcast: filters all types of multicast packets, including known. Not supported when resource-mode is on. • bcast: filters all broadcast packets • unknown-l2-mcast: filters all unknown L2 multicast packets — L2-only devices: includes all unknown multicast frames, including unknown IP multicast frames — L3-capable devices: includes all unknown multicast frames except those with IP multicast addresses • unknown-ip-mcast: filters all unknown IP multicast packets. Not supported when resource-mode is off. This option is not available on L2-only devices (such as 3911) — In case of MPLS configurations on a port or a MPLS virtual-switch, IP multicast frames are treated as L2 multicast frames. — Cannot be configured on a filter which has VPLS VS added as classifier. • unknown-ucast: filters all unknown unicast packets By default, the filtered traffic types are Layer 2 Broadcast traffic (bcast) and Layer 2 Unknown Multicast traffic (l2-mcast). ***Note:*** Downgrading to a software release prior to 6.5 results in a loss of any containment classification settings. Filtered traffic types revert to Layer 2 Broadcast traffic (bcast) and Layer 2 Unknown Multicast traffic (l2-mcast).

—end—

Copyright© 2022 Ciena® Corporation

Procedure 12-3
Setting broadcast containment configuration

Set broadcast containment configuration to

- enable or disable the resource mode
- instruct the system whether to ignore R-APS messages when filtering PDUs

Step	Action
1	Set the broadcast containment configuration: `broadcast-containment set {resource-mode <on|off>} {ignore-raps-msgs <on|off>}` where resource-mode <on | off> — determines whether the resource usage mode is on or off. ignore-raps-msgs <on | off> — determines whether R-APS messages are ignored when filtering PDUs. **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 12-4
Enabling and disabling broadcast containment

Enable broadcast containment to use the broadcast containment feature. Disable broadcast containment when you no longer require the broadcast containment feature.

By default, broadcast containment is disabled.

Step	Action
To enable broadcast containment	
1	Enable broadcast containment: `broadcast-containment enable`
To disable broadcast containment	
2	Disable broadcast containment: `broadcast-containment disable`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 12-5
Deleting a broadcast containment filter

Delete a broadcast containment filter.

Note: Use caution when deleting broadcast containment filters. The system software does not prevent you from deleting an active broadcast containment filter and automatically removes the associated ports from the filter.

Step	Action
1	Delete a broadcast containment filter: `broadcast-containment delete filter <filter>` where filter <filter> is the filter to be deleted. **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 12-6
Adding ports to a broadcast containment filter

After creating a broadcast containment filter, you can add additional ports to it.

Note: Multiple ports can be added to a filter, but only one filter may be applied to a specific port at a time. Also, when multiple ports are added to a filter, the bandwidth for the filter is shared among the ports.

Step	Action
1	Add a port to a broadcast containment filter: `broadcast-containment add filter <filter> {port <Port list>}` where filter <filter> — Is the filter that you want to add ports to. port <Port list> — is the list of ports to add to a filter. **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 12-7
Removing a port from a broadcast containment filter

Remove a port from a broadcast containment filter when you no longer want the broadcast containment filter to be applied to the port.

Step	Action
1	Remove a port from a broadcast containment filter: `broadcast-containment remove filter <filter> port <Port list>` where filter <filter> — Is the filter that you want to remove ports from. port <Port list> — is the list of ports to remove from a filter. **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 12-8
Displaying broadcast containment filters

Display broadcast containment filters when you want to confirm the configuration and confirm that the configuration is working as planned.

Step	Action

To display broadcast containment status and a summary of broadcast containment filters

1 Display broadcast containment status and a summary of broadcast containment filters:

```
broadcast-containment show
```

To display details about a specific broadcast containment filter

2 Display details about a specific broadcast containment filter:

```
broadcast-containment show [filter <filter>] {statistics}
```

where

filter <filter> is the filter.

To display details about a specific broadcast containment port

3 Display details about a specific broadcast containment filter:

```
broadcast-containment show [port <port>]
```

where

port <port> is the port list.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example displays broadcast containment status and a summary of broadcast containment filters.

```
broadcast-containment show

> broadcast-containment show

+-------- BROADCAST CONTAINMENT GLOBALS ---------+
|         Parameter       |        Value         |
+-------------------------+----------------------+
| Admin Status            | Disabled             |
| Resource Mode           | On                   |
| Ignore R-APS Messages   | Off                  |
| Global Dropped Bytes    | 0                    |
| Global Dropped Packets  | 0                    |
+-------------------------+----------------------+


+-------- BROADCAST CONTAINMENT FILTERS ---------+
|          Filter         |           |          |
| Name            | ID    | Kbps      | Dropping |
+-----------------+-------+-----------+----------+
| denial          | 1     | 2048      | No       |
+-----------------+-------+-----------+----------
```


Copyright© 2022 Ciena® Corporation

This example displays details about a specific broadcast containment filter with a port as a member in resource mode. (This is not available on 390x platforms on which MPLS is not supported.)

```
> broadcast-containment show filter filter1

+---------- Broadcast Containment Filter ----------+
| Parameter                  | Value               |
+----------------------------+---------------------+
| Filter Name                | filter1             |
| Index                      | 1                   |
| Kilobits Per Second (Kbps) | 0                   |
| Broadcast                  | Off                 |
| Unknown Unicast            | On                  |
| Unknown L2 Multicast       | Off                 |
| Unknown IP Multicast       | Off                 |
+----------------------------+---------------------+
+------------------ ATTACHMENTS ------------------+
|                   CLASSIFIERS                   |
| Port                    | Virtual-Switch        |
+-------------------------+-----------------------+
| 4                       | -                     |
+-------------------------+-----------------------+
| 5                       | -                     |
+-------------------------+-----------------------+
```


Copyright© 2022 Ciena® Corporation

This example displays details about a specific broadcast containment filter with a port as a member in non-resource mode. (On 390x output is the same in resource mode and in non-resource mode.)

```
> broadcast-containment show filter filter3

+---------- Broadcast Containment Filter ----------+
| Parameter                  | Value                |
+----------------------------+----------------------+
| Filter Name                | filter3              |
| Index                      | 3                    |
| Kilobits Per Second (Kbps) | 64                   |
| Broadcast                  | Off                  |
| Unknown Unicast            | On                   |
| All Multicast              | Off                  |
+----------------------------+----------------------+
+---------- ATTACHMENTS -----------+
|           CLASSIFIERS            |
| Port                             |
+----------------------------------+
| 3                                |
+----------------------------------+
```


Copyright© 2022 Ciena® Corporation

This example displays details about a specific broadcast containment filter with per port and per VS members attached to it in resource mode. (This is not available on 390x platforms on which MPLS is not supported.)

```
> broadcast-containment show filter filter2

+---------- Broadcast Containment Filter ----------+
| Parameter                  | Value               |
+----------------------------+---------------------+
| Filter Name                | filter2             |
| Index                      | 2                   |
| Kilobits Per Second (Kbps) | 0                   |
| Broadcast                  | Off                 |
| Unknown Unicast            | On                  |
| Unknown L2 Multicast       | Off                 |
| Unknown IP Multicast       | Off                 |
+----------------------------+---------------------+
+------------------ ATTACHMENTS ------------------+
|                   CLASSIFIERS                   |
| Port                    | Virtual-Switch        |
+-------------------------+-----------------------+
| 1                       | VS1                   |
+-------------------------+-----------------------+
| 2                       | VS2                   |
+-------------------------+-----------------------+
```


Copyright© 2022 Ciena® Corporation

This example displays statistics for a specific broadcast containment filter . (This is not available on 390x platforms on which MPLS is not supported.)

```
> broadcast-containment show filter filter2

+---------- Broadcast Containment Filter ----------+
| Parameter                  | Value                |
+----------------------------+----------------------+
| Filter Name                | filter2              |
| Index                      | 2                    |
| Kilobits Per Second (Kbps) | 0                    |
| Broadcast                  | Off                  |
| Unknown Unicast            | On                   |
| Unknown L2 Multicast       | Off                  |
| Unknown IP Multicast       | Off                  |
| Dropped Bytes              | 0                    |
| Dropped Packets            | 0                    |
+----------------------------+----------------------+
+------------------ ATTACHMENTS ------------------+
|                   CLASSIFIERS                   |
| Port                  | Virtual-Switch         |
+-----------------------+------------------------+
| 1                     | VS1                    |
+-----------------------+------------------------+
| 2                     | VS2                    |
+-----------------------+------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 12-9
Displaying a broadcast containment filter to which per port per VS membership is attached

Use this procedure to display a broadcast containment filter to which per port per VS membership is attached.

Step	Action
1	Display a broadcast containment filter to which per port per VS membership is attached: `broadcast-containment show virtual-switch <name> port <name>` where virtual-switch <name> — specifies the virtual switch for per port per VS membership port <name> — specifies the port for per port per VS membership **—end—**

Example

This example displays broadcast containment status and a summary of broadcast containment filters.

```
> broadcast-containment show virtual-switch VS1 port 1
+------ VIRTUAL-SWITCH VS1 PORT 1 ------+
| Parameter                   | Value    |
+-----------------------------+----------+
| Filter Name                 | filter2  |
| Kilobits Per Second (Kbps)  | 0        |
| Broadcast                   | Off      |
| Unknown Unicast             | On       |
| Unknown L2 Multicast        | Off      |
| Unknown IP Multicast        | Off      |
+-----------------------------+----------+
```


Copyright© 2022 Ciena® Corporation

Procedure 12-10
Displaying broadcast containment configuration in the configuration file

Display broadcast containment configuration in the configuration file when you want to

- confirm the configuration
- copy the configuration to an file

Broadcast containment configuration is saved to the configuration file in the BCAST FILTER CONFIG: section.

Step	Action
1	Display broadcast containment configuration in the configuration file: `configuration show` **—end—**

Example

This example shows sample output for the configuration show command.

```
> configuration show
...
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
! BCAST FILTER CONFIG:
!
broadcast-containment create filter denial kbps 2048
!
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
```


Copyright© 2022 Ciena® Corporation

Procedure 12-11
Displaying all broadcast containment resources

Use this procedure to display the resources being used by broadcast containment.

Step	Action
1	Show resources being used by broadcast containment. `broadcast-containment show resources` **—end—**

Example

```
> broadcast-containment show resources

+-- BROADCAST CONTAINMENT RESOURCE USAGE--+
| Parameter                |     Count      |
|                          | Total | In Use |
+--------------------------+-------+--------+
| Filters                  | 256   | 1      |
| Member Resource Pool     | 1459  | 2      |
+--------------------------+-------+--------+
```


Copyright© 2022 Ciena® Corporation

Procedure 12-12
Configuring broadcast containment filters on a port and classifier combination of port and virtual switch

This example shows how to rate-limit traffic on a port only and on classifier combination of per port per virtual switch basis.

Step	Action

To configure virtual switches to be ready for broadcast containment

1 Create the first VPLS virtual switch (VS):

```
virtual-switch ethernet create vplsvs1 mode vpls
```

2 Create the second VPLS VS:

```
virtual-switch ethernet create vplsvs2 mode vpls
```

3 Add the first port to vplsvs1:

```
virtual-switch ethernet add vs vplsvs1 port 1 vlan 100
```

4 Add the first port to vplsvs2

```
virtual-switch ethernet add vs vplsvs2 port 1 vlan 200
```

To configure to rate-limit traffic on a port

5 Create a filter with containment classification and rate limit:

```
broadcast-containment create filter uc1 containment-
classification unknown-ucast kbps 10
```

6 Add a port to the filter:

```
broadcast-containment add filter uc1 port 1
```

7 Enable broadcast containment:

```
broadcast-containment enable
```

To configure to rate-limit traffic on a per port per virtual switch basis

8 Create a filter with containment classification and rate limit:

```
broadcast-containment create filter uc2 containment-
classification unknown-ucast kbps 100
```

9 Add a port and a virtual switch to the filter:

```
broadcast-containment add filter uc2 vs vplsvs1 port 1*
```

10 Enable broadcast containment:

```
broadcast-containment enable
```

—end—

Note 1: In the above configuration examples, members added to both filters have one classifier in common, which is port 1.

Copyright© 2022 Ciena® Corporation

Note 2: Unknown-ucast traffic qualifying on port 1 and virtual-switch “vplsvs1” are rate-limited on the basis of broadcast filter “uc2” as the classifier combination of port and virtual switch takes precedence over the single classifier of port.

Note 3: Unknown-ucast traffic qualifying on port 1 and virtual-switch “vplsvs2” are rate-limited on the basis of broadcast filter “uc1”.

Copyright© 2022 Ciena® Corporation

Procedure 12-13
Clearing broadcast containment statistics

Use this procedure to clear the statistics for all broadcast containment filters, or to clear the statistics for a specific broadcast containment filter.

Step	Action
To clear the statistics for all broadcast containment filters	
1	Clear the statistics for all broadcast containment filters: `broadcast-containment clear statistics`
To clear the statistics for a specific broadcast containment filter	
2	Clear the statistics for a specific broadcast containment filter: `broadcast-containment clear filter <filter> statistics`

—end—

Copyright© 2022 Ciena® Corporation

Copyright© 2022 Ciena® Corporation

Multi-chassis link aggregation

A LAG uses multiple Ethernet links to increase bandwidth and provide link-level redundancy. A multi-chassis (MC) LAG is a LAG with ports that terminate on separate chassis to provide node-level redundancy, along with the link-level redundancy provided on the LAG. MC-LAG provides equipment protection in dual-home topologies by maintaining a forwarding path for all services if one node fails. It models two or more network elements as one single node, eliminating loops among the interconnection nodes. Two or more nodes can share a common LAG endpoint. For more information about LAGs, see the *Link Aggregation* chapter in *39XX/51XX Switches and Platforms Base Configuration.*

MC-LAG is based on IEEE 802.1AX. This standard uses Link Aggregation Groups with Link Aggregation Control Protocol (LACP). The switches in the redundancy group communicate with each other so that redundancy group operates properly. An inter-chassis link (ICL) protocol runs among the provider edge (PE) devices in the redundancy group to synchronize configuration and/or running data.

The link between systems operating LACP is assigned a unique priority in this order: System Priority, System ID, Port Priority and Port Number. The highest priority has the lowest numerical value. A port that is attached to the highest priority link is available for active use. The port is selected for active use if the preceding higher priority selection can be maintained. If it cannot be maintained, the port is selected as a standby.

The switch with the higher system priority selects the standby port. The switch can be a redundancy group or dual-homed device. The switch selects the standby link based on the port's priority with the higher priority port becoming the active port. The higher system priority switch selects the active port based on local priority. A lower system switch selects the active port based on remote port priority.

MC-LAG supports the triangle ICL topology which consists of one active and one standby chassis in the redundancy group. The Configuring MC-LAG figure shows MC-LAG configuration where the active and standby links are connected to the same peer switch. For a shared co-located dedicated

Copyright© 2022 Ciena® Corporation

interconnect, the PE devices in a redundancy group are co-located in the same physical location. The interconnect for ICL is provided through the core network to which the PE devices are connected. For a geo-redundant shared interconnect, the PE devices in a redundancy group are located in different physical locations. The interconnect for ICL is provided over the network in which the PE devices are connected.

Note: A 1+1 MC-LAG is the only configuration supported.

Figure 13-1
Configuring MC-LAG

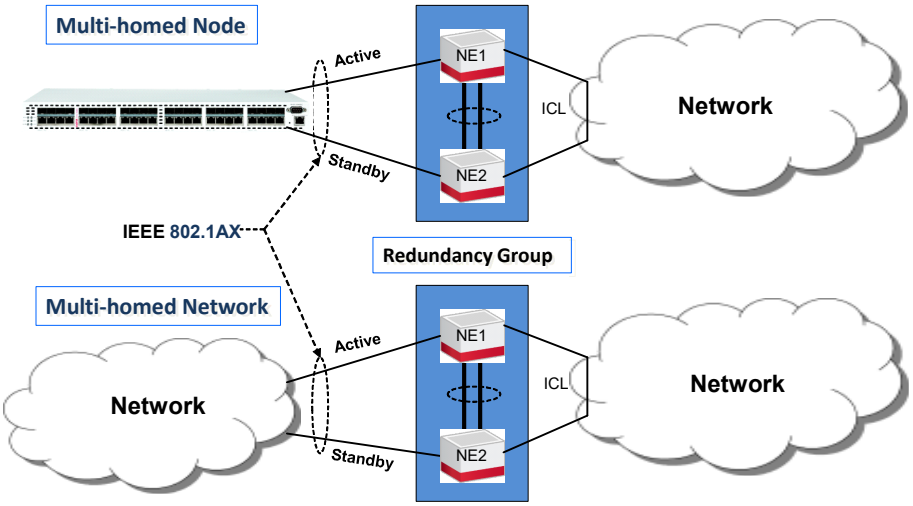

Hashing

The hashing mechanism common across all LAGs also works on MC-LAG configurations.The hashing distribution is not apparent because MC-LAG only supports 1+1 configurations. The hashing algorithm is not applicable to control frames exchanged over an inter-chassis link. For more information about hashing, see the *Link Aggregation* chapter in *39XX/51XX Switches and Platforms Base Configuration.*

Monitoring inter-chassis link connections

MC-LAG detects redundancy group failures and recoveries in seconds, but Connectivity Fault Management (CFM) and Label Switch Path (LSP) Bidirectional Forwarding Detection (BFD) provide faster fault detection and switchover. Without CFM and BFD, MC-LAG relies on the Inter-Chassis Protocol (ICP) keep-alive message which takes longer to detect failed PE devices.

Copyright© 2022 Ciena® Corporation

CFM needs to be configured at a Maintenance End Point (MEP) on each east and west ring span of the ICL VLAN or Virtual Switch (VS). BFD must be configured on the primary and backup tunnel protection on the ICL VS. The CFM and BFD service names are provided to MC-LAG during ICL configuration. The MC-LAG module consolidates configuration, administration and alarm states of the CFM services to monitor the peer node. Using CFM Down MEPs also supports hardware accelerated CFM.

Detecting failures

An active link failure detected by an active node triggers the link switchover. The active node sends an ICL protocol data unit (PDU) to inform the standby PE to switchover after the link failure.

If BFD or CFM is not configured, MC-LAG detects node failure based on ICL heart beat time out. ICL periodically sends keep-alive PDUs to the peer node. If no PDUs are received from the peer during the heartbeat interval, then the peer is deemed unreachable.

On the standby switch, CFM or BFD detects the failure and informs MC-LAG to act as a new active node. The link on the new active mode is then activated.

> ***Note:*** The Rx Timeout Multiplier for an ICL is not configurable on SAOS 6.x platforms (as it is on SAOS 8.x platforms), but has a default value of 3.

MC-LAG interoperability between SAOS 6.x and 8.x

The user can create 1+1 MC-LAG using SAOS 6.x and 8.x nodes as redundancy group (RG) members. The admin-key value has increased its range from 2049-4095 to 2049-8191 to accommodate this interoperability.

Procedures

For procedures related to multi-chassis link aggregation, see:

- “Configuring a multi-chassis link aggregation group” on page 13-4
- “Adding a port to a multi-chassis link aggregation group” on page 13-6
- “Adding a protection port to a multi-chassis link aggregation group” on page 13-7
- “Forcing a switchover on a redundancy group node” on page 13-8
- “Configuring an inter-chassis link” on page 13-9
- “Displaying an inter-chassis link” on page 13-10
- “Displaying a summary of MC-LAGs” on page 13-11

See *39xx/51xx Switches and Platforms Command Reference* for additional MC-LAG commands.

Copyright© 2022 Ciena® Corporation

Procedure 13-1
Configuring a multi-chassis link aggregation group

Configure an MC-LAG for node level redundancy. These restrictions apply:

- Only one port can be added to an MC-LAG, either a primary port or a protection port.
- A LAG cannot be converted to MC-LAG using the “aggregation set agg <agg-name> multi-chassis” command if
  - Agg is in manual mode
  - Agg has member ports already added to it

The MC-LAG mode has to be LACP. MC-LAG does not support manual mode.

Step	Action
1	Create a link aggregation group. `aggregation create agg <agg-name>` where agg-name — is the name of the LAG.
2	Create or configure the MC-LAG: `aggregation set agg <agg-name> multi-chassis hash {mac-address-based | ip-addr-based | enhanced} revert-delay {0-720,000 ms} system-priority {0-65535} system-mac {mac-address} admin-key {2049-8191} protection-mode {proprietary | standard} revert-protection {on | off} [icl {icl-name}]` where agg-name> — is the name of the MC-LAG. multi-chassis — converts a LAG to a MC-LAG. hash {mac-address-based | ip-addr-based | enhanced} — hashing mode for known unicast frames. The default value is mac-address-based. revert-delay {0-720,000 ms} — sets the reversion delay from 0 to 720,000 ms The default is 300,000 ms for MC-LAG. system-priority {0-65535} — sets the LAG redundancy group priority. This value must be the same on both redundancy group nodes.

Copyright© 2022 Ciena® Corporation

where

system-mac {mac-address}	sets the user-defined redundancy group MAC address. This value must be the same on both redundancy group nodes. ***Note:*** The system-priority and the system-mac form the redundancy group's system-id. The system-id value must be the same on both redundancy group nodes.
admin-key {2049-8191}	sets the LAG admin key. This value must be the same on both redundancy group nodes.
protection-mode {proprietary | standard}	sets the aggregation protection mode
revert-protection {on | off}	sets revert protection on or off
icl-name	optional attribute that indicates the inter-chassis link (ICL) that the redundancy group uses

3 Verify the MC-LAG configuration:

```
aggregation show agg <agg-name>
```

where

agg-name	is the name of the MC-LAG.

—end—

Example

This example shows the creation and configuration of an MC-LAG.

```
aggregation create agg MC-LAG1

aggregation set agg MC-LAG1 multi-chassis hash mac-addr-based
revert-delay 65000 system-priority 46000

system-mac 00:E2:F2:F1:22:01 icl ICL2 admin-key 4000
```


Copyright© 2022 Ciena® Corporation

Procedure 13-2
Adding a port to a multi-chassis link aggregation group

Add a port to the MC-LAG redundancy group. Adding a port as a primary port to the redundancy group makes it the primary redundancy group node.

Note: An MC-LAG can contain only one primary port.

Step	Action
1	Add a port to an MC-LAG: `aggregation add agg <agg-name> port <port-num>` where agg-name — is the name of the MC-LAG port-num — is the port number you are adding to the MC-LAG. **—end—**

Example

This example shows how to add a port to the MC-LAG.

```
aggregation add agg MC-LAG1 port 1
```


Copyright© 2022 Ciena® Corporation

Procedure 13-3
Adding a protection port to a multi-chassis link aggregation group

Add a protection port to the MC-LAG redundancy group. Adding a port as the protection port to the redundancy group makes it the standby redundancy group node.

Note: An MC-LAG can contain only one protection port.

Step	Action
1	Add a protection port to a MC-LAG: `aggregation add agg <agg-name> protection-port <port-num>` where agg-name — is the name of the MC-LAG. port-num — is the port number you are adding to the MC-LAG. **—end—**

Example

This example shows how to add a protection port to the MC-LAG.

```
aggregation add agg MC-LAG1 protection-port 1
```


Copyright© 2022 Ciena® Corporation

Procedure 13-4
Forcing a switchover on a redundancy group node

You can force a switchover from an active redundancy group node to an inactive redundancy node on a MC-LAG. This action is displayed on both the active and inactive nodes, but only takes effect on the active node.

A reversion may occur after switchover if the command is executed on the primary node and reversion is on, meaning that reversion is enabled on backup.

Step	Action
1	Force a switchover from an active redundancy group node on a MC-LAG: `aggregation switchover agg <agg-name>` where agg-name is the name of the MC-LAG. **—end—**

Example

This example shows how to force a switchover.

```
aggregation switchover agg MC-LAG1
```


Copyright© 2022 Ciena® Corporation

Procedure 13-5
Configuring an inter-chassis link

Create an inter-chassis link (ICL) on a VLAN or VS.

Step	Action
1	Create an ICL on the VLAN or VS `icl create link <icl-name> vlan|vs <vlan-name or vlan-id | vs-name> remote-mac-address <mac-address> heartbeat-interval <1-30s> cfm-service-primary <service-name> cfm-service-secondary <service-name>`

where

icl-name	is the inter-chassis link name
vlan-name or vlan-id | vs-name	is the VLAN or VS that the ICL uses
mac-address	is the ICL peer's chassis MAC address
heartbeat-interval <1-30s>	is an optional parameter that specifies the heartbeat interval. The default is 5 seconds
cfm-service-primary <service-name	is the name of the CFM service monitoring the primary ICL path
cfm-service-secondary <service-name>	is the name of the CFM service monitoring the secondary ICL path

—end—

Example

This example shows how to configure an ICL.

```
icl create link ICL1 vlan vlan1 remote-mac-address 10:E2:e3:F1:44:01
heartbeat-interval 20 cfm-service-primary cfmservice1 cfm-service-secondary
cfmservice2
```


Copyright© 2022 Ciena® Corporation

Procedure 13-6
Displaying an inter-chassis link

Display an inter-chassis iink (ICL) configured on the system.

Step	Action
1	Display an inter-chassis link: `icl show link <icl-name>` where icl-name is the inter-chassis link name **—end—**

Example

```
icl show link myIcl
+------------------------------- ICL myIcl Info -------------------------------+
|          Parameter          |            Value                                |
+-----------------------------+-------------------------------------------------+
| ICL Name                    | myIcl                                           |
| ICL ID                      | 1                                               |
| Admin State                 | Enabled                                         |
| Oper  State                 | Enabled                                         |
| ICL Status                  | Active                                          |
| Remote Mac Address          | 00:23:8A:A4:9A:B0                               |
| Vlan/VS Using               | 99                                              |
| CFM Service Primary         |                                                 |
| CFM Service Secondary       |                                                 |
| ICL Faults                  |                                                 |
|    CFM Fault                | No                                              |
|    VPLS Fault               | No                                              |
|    Rx Timeout               | No                                              |
|    Interval Mismatch        | No                                              |
| Heartbeat Interval          | 1s                                              |
| Up Time                     | 5s                                              |
| Total Down Time             | 10s                                             |
| Rx Frames                   | 16                                              |
| Tx Frames                   | 34                                              |
| Rx Unknown Frames           | 0                                               |
| Rx Heartbeat Frames         | 7                                               |
| Tx Heartbeat Frames         | 21                                              |
| Rx Config Mismatch Frames   | 1                                               |
| Tx Failed Frames            | 0                                               |
| Number of Failures          | 1                                               |
+-----------------------------+-------------------------------------------------+
```

Note: The Rx Timeout Multiplier for the ICL does not appear beneath Heartbeat Interval in the display (as it does on SAOS 8.x platforms) because it is not configurable on SAOS 6.x platforms. It has a default value of 3.

Copyright© 2022 Ciena® Corporation

Procedure 13-7
Displaying a summary of MC-LAGs

Display a summary of MC-LAGs configured on the system.

Step	Action
1	Display a summary of MC-LAGs: `aggregation show multi-chassis` **—end—**

Copyright© 2022 Ciena® Corporation

Copyright© 2022 Ciena® Corporation

CFM-based link aggregation

In general, LAG switchover is based on either a port down event or LACP timeout.

In the case where switchover occurs due to a port down event, then switchover may occur within 50ms, which is not guaranteed because this event may have inherent latency (due to a delay introduced by a software-based link-scan mechanism). Features like port holdoff can also delay the detection of a link down event.

In the case where a LAG is formed by an indirectly connected node exchanging LACPDU through the L2CFT tunnel, then switchover time is very high as it is triggered by LACP timeout, and LACP timeout is equal to 3xTimeout. Timeout can be set to either 1 second or 30 seconds. So, at best, switchover time (caused at remote end of LACP partner node not sending LACPDU) can be 3 seconds.

As a solution to these problems, a new “Connectivity Fault Management (CFM) over LAG members” feature is added in SAOS 6.18. This feature allows the user to create a down MEP on the physical member of a LAG, and the CCMs of this MEP are exchanged only on that physical member port of the LAG. Whenever there is a connectivity issue at the remote or intermediate node, it may be detected by either or one of the LACP partners but not both of them. But MEPs at both sides can detect a CFM fault and the service within which the fault occurs then reports the fault to the associated LAG member.

As of SAOS 6.18, MEPs on LAG member ports are supported only on accelerated CFM services. This functionality is supported on all SAOS 6.18 active platforms.

As of SAOS 6.21, CFM monitoring of both active LAG member ports and protected LAG member ports is supported to allow for faster switching between ring paths.

The user can now associate a LAG member with a CFM service, as shown in step 9 of the procedure “Configuring CFM service fault-based LAG switchover” on page 14-4.

Copyright© 2022 Ciena® Corporation

CFM service fault-based LAG switchover is supported for:

- LAG in “manual” mode
- LAG in “LACP” mode
- Active/Active LAG
  - — The LAG switches over traffic from the failed member port to other LAG members when the fault is reported by the CFM service associated with that port.
  - — The LAG moves the port back into distribution when it gets the “fault clear” event from the CFM service associated with that port.
- Active/Standby LAG
  - — The LAG moves the primary port out of distribution and fills the vacancy by moving a protection port (if available) into distribution from a standby state when the fault is seen on the CFM service associated with the primary port.
  - — When it gets the “fault clear” event from the CFM associated with the primary port of the LAG, the LAG moves the port into:
    - – Standby state, if there is no vacancy in the LAG.
    - – Distribution, if there is a vacancy in the LAG.
- Reversion after fault clear event—The LAG supports revertive and non-revertive modes when the CFM service associated with the primary member of the LAG reports the “fault clear” event, as mentioned here:
  - — When reversion is enabled, then after the reversion time expires on the primary port, it becomes an active member of the LAG by bumping a protection port out of distribution.
  - — When reversion is disabled, then the primary port remains in a standby state.
- Any fault in the CFM service triggers the LAG switchover when the user associates the CFM service with the member of the LAG.
- LAG does not switchover the traffic and ignores the fault event received for any given LAG member if the CFM service that reported the fault was not associated with the port.

  Note: This is intended to avoid false switchovers due to initial faults like “IS” seen during the MEP discovery phase. It is recommended to associate the CFM service with the LAG member port being monitored only when the CFM service has come up properly and all faults (especially “IS”) are cleared.

Copyright© 2022 Ciena® Corporation

- When the user removes a LAG member port from the LAG, then:
  - — If the LAG member port being removed is associated with a CFM service, then this CFM service is dissociated from that physical port.
  - — If a MEP exists on LAG member port, then that MEP is also deleted.
- The LAG ignores the “fault clear” event from MEP running on the LAG member port in the following scenarios:
  - — When no CFM service is associated with the LAG member port.
  - — When the CFM service associated with the member port is not the same from which the event was received.

Copyright© 2022 Ciena® Corporation

Procedure 14-1
Configuring CFM service fault-based LAG switchover

Associate a LAG member with the CFM service used to monitor the health of the LAG member port. The same CFM service cannot be associated with more than one LAG member port.

A fault on the CFM service is used as a trigger to remove the associated member port out of distribution.

> ***Note:*** Only a primary LAG member can be associated with a CFM service.

You can also dissociate a LAG member port from a CFM service.

Step	Action

To configure CFM service fault-based LAG switchover

1 Create the LAG and add member ports:

a. Create a LAG on the device:

```
aggregation create agg <String[31]>
```

where

agg <String[31]>	is the aggregation port name.

b. Add ports to the LAG:

```
aggregation add agg <agg> port <Physical port list>
protection-port <Physical port list>
```

where

agg <agg>	is the aggregation port.
port <Physical port list>	is the list of ports to add to the aggregation.
protection-port <Physical port list>	is the list of protection ports to add to the aggregation.

Copyright© 2022 Ciena® Corporation

2 Create VLAN(s) and add LAG to them:

a. Create a VLAN for the LAG

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

where

vlan <vlan>	is the VLAN ID.
name <String[31]>	is the name of the VLAN.

b. Add the aggregation to the VLAN

```
> vlan add vlan <vlan> port <Port list>
```

where

vlan <vlan>	is the VLAN ID.
port <Port list>	is the list of ports to add to the VLAN.

3 Enable CFM:

```
> cfm enable
```

4 Create an accelerated CFM service on the VLAN:

```
> cfm service create vlan <vlan> name <String[45]>
accelerate
```

5 Assign an MD level of 0 to the CFM service for optimal performance:

```
> cfm service set service <service-name> md md0
```

6 Assign an alarm priority of 3, an alarm time of 0, and a CCM interval of 3.3 or 10 ms to the CFM service for optimal performance:

```
> cfm service set service <service-name> alarm-priority 3
alarm-time 0 ccm-interval 3.3
```

7 While creating MEP, provide an additional “member-port’ parameter. This is applicable when down MEP is created on the LAG:

```
> cfm mep create service <service-name> port <agg-
name> vlan <vlan-id> type down mepid <mep-id> member-port
<port-number>
```

8 Enable the CFM service:

```
> cfm service enable service <service-name>
```

9 Associate a LAG member with a CFM service:

```
> aggregation set port <port-no> cfm-service <service-
name>
```

To dissociate a LAG member from a CFM service

10 When needed, dissociate a LAG member from a CFM service:

```
> aggregation unset port <port-no> cfm-service
```

—end—

Copyright© 2022 Ciena® Corporation

Copyright© 2022 Ciena® Corporation

Port state mirror group

This section describes port configuration for fault management and provides procedures for configuring ports for fault management.

Topics are

- “Port mirroring” on page 15-1
- “Port state mirroring” on page 15-3
- “Limitations” on page 15-6

Port mirroring

Port mirroring is where another port is configured to “mirror” the traffic of a specified port to monitor ingress and egress traffic. By attaching a protocol analyzer to the port, you can monitor and analyze traffic. Port mirroring actions can only be configured on physical ports.

Monitoring traffic on a port is a two-step process. First, port mirroring must be enabled on the port that is to be the mirror port. Next, direct the input traffic, output traffic, or both to the mirror port (the port to which the analyzer is attached).

Any port can be designated as a mirror port. The maximum number of mirror-ports that can be configured is 2. The ingress/egress traffic from multiple ports can be mapped to a single mirror port. All egress-traffic, from any source port, must be mapped to the same mirror-port, while ingress traffic from some source ports can be mapped to one of the two possible mirror-ports, and ingress traffic from other source ports can be mapped to the second mirror-port.

> ***Note 1:*** TWAMP frames are excluded from port mirroring if those frames are processed by hardware acceleration.
>
> ***Note 2:*** F-CoS and R-CoS attributes are not applied to the mirroring port.

Port mirroring configuration persists after a reboot if the device configuration has been saved. Note that a port defined as a mirror port continues to transmit and receive normally based on its configuration, such as its VLAN membership.

Copyright© 2022 Ciena® Corporation

Mirroring can be performed in one of two modes:

- Standard port mirroring, which causes mirrored packets to be transmitted out of the mirror-port without any frame modifications.
- Encapsulation mode, which causes an extra vlan-tag to be pushed onto the mirrored frame as it egresses the mirror-port.

Mirror frame encapsulation

There is no restriction on the VID used in the vlan-tag that can be added to mirrored packets in encapsulation mode. The tag is added only as the mirrored-packet exits the mirror-port: no switching is done on this tag by the node generating the mirrored traffic. The VID used in the tag need not belong to a VLAN that is extant on the device generating the mirrored traffic. If the VLAN is extant on that device, the mirror-port does not need to be a member of that VLAN.

The packet being mirrored can be of any format. The mirroring encapsulation functionality is ignorant of the nature of the mirrored traffic. It simply adds an extra tag to the mirrored packet, without any knowledge of the existing content of that packet. Thus, with mirroring encapsulation employed, untagged source packets are mirrored single-tagged, single-tagged source packets are mirrored double-tagged, double-tagged packets are mirrored triple-tagged, and so on.

> ***Note:*** To avoid frame duplication, QoS shaping bottlenecks etc in standard mode, it is recommended that the destination mirror port be a stand-alone port used for diagnostic purposes. Normal traffic can be affected if the destination mirror port is included in the path of normal virtual-switch interface configurations.

Egress port mirroring behavior

Depending on the SAOS device, egress port mirroring has one of these types:

- traditional egress mirroring
- true egress mirroring

Traditional egress mirroring frames destined for a specific egress port, with the mirrored frame transmitted by the mirror port resembling the original frame before egress frame modifications occur.

Examples of egress frame modifications that do not appear in traditional egress mirrored frames include

- egress TPID setting
- egress untagged VLAN setting
- VLAN translation on the egress port

Copyright© 2022 Ciena® Corporation

Traditional mirrored frames may contain transport encapsulations that were received or temporarily pushed on by the device. This may occur where egress mirroring mirrors these types of traffic:

- egressing a VS member on a Q-in-Q virtual switch with or without a virtual circuit assign
- egressing a VS member of a MPLS or BPT virtual switch
- egressing a sub-port interface

True egress mirroring mirrors frames destined for a specific egress port, where the mirrored frames transmitted by the mirror port resembling as it was transmitted out the specific egress port, including all egress modifications to the frame.

This table shows which devices are supported on the two types of egress mirroring.

Table 15-1
Egress mirroring

Egress mirroring type	Devices supported
Traditional egress mirroring	• 3903 • 3903x • 3904 • 3905 • 3906
True egress mirroring	• 3924 • 3926 • 3928 • 3942 • 5142 • 5160

Procedures for port mirroring are:

- “Configuring port mirroring” on page 15-8
- “Configuring encapsulated port mirroring” on page 15-9

Port state mirroring

Port state mirroring (sometimes referred to as link loss forwarding or link state tracking) provides link redundancy by associating the link state of one or more source (uplink) ports with one or more downstream (destination) ports on the switch. If the link is lost on a source port, all of the other destination ports

Copyright© 2022 Ciena® Corporation

associated with it are automatically placed in a disabled state as well. When used in conjunction with RSTP, port state mirroring can be leveraged to force the use of a secondary link.

Port mirroring also allows encapsulation of mirrored packets with an outer vlan-tag, as those mirrored packets egress the mirror-port.

Allowing the port state of source ports to directly control the port state of destination ports provides an easy way to shape traffic patterns. Port state mirroring also provides another way to detect link health that otherwise might not be detected.

If port state mirroring is not configured and a link is lost (a cable is disconnected, the connected switch or router goes down, etc.) the link state of the associated interfaces is not affected and a failover to a secondary interface is not initiated.

Uni-directional port state mirroring

Port state mirroring allows ports to be bundled together in a related group referred to as a port state mirror group (PSMG). In the group, the link state of the destination ports depends on the link state of the source port(s) in the PSMG. With a single source port, if the source port becomes disabled, all of the destination ports in the group are forced into a disabled state as well. Likewise, when the source port becomes available again, the destination ports transition back to an enabled state (depend on their own availability).

In a PSMG that has more than one source port, all source ports must be disabled in order for the destination ports to be forced into a disabled state. For example, if a PSMG has ports 1 and 2 configured as source ports, both ports must be down to force the destination ports into a disabled state. This also means that only one source port is required to be up to transition all destination ports back to the enabled state.

In Figure 15-1, a PSMG is configured on Device 1 and another PSMG on Device 2. On Device 1, the downstream (destination) ports 1, 2, 5, and 6 are defined in PSMG 1 with uplink ports 3 and 4 acting as the source ports. For Device 2, downstream ports 7, 8, 11, and 12 are defined in PSMG 2 with uplink ports 9 and 10 acting as the source ports. Ports 1 and 2 on Device 1 serve as the primary connections for Server 1 and Server 2, whereas port 7 and 8 on Device 2 are the secondary links for Server 1 and Server 2. Likewise, the primary links for Server 3 and Server 4 are through port 11 and 12 on Device 2 with secondary links to Device 1 through ports 5 and 6.

Copyright© 2022 Ciena® Corporation

When the uplinks become unavailable, the servers are switched over to the secondary links and the primary links are disabled. Primary links are represented as solid lines, whereas secondary links are represented as dashed lines.

This figure shows port state mirroring configuration.

Figure 15-1
Port state mirroring configuration

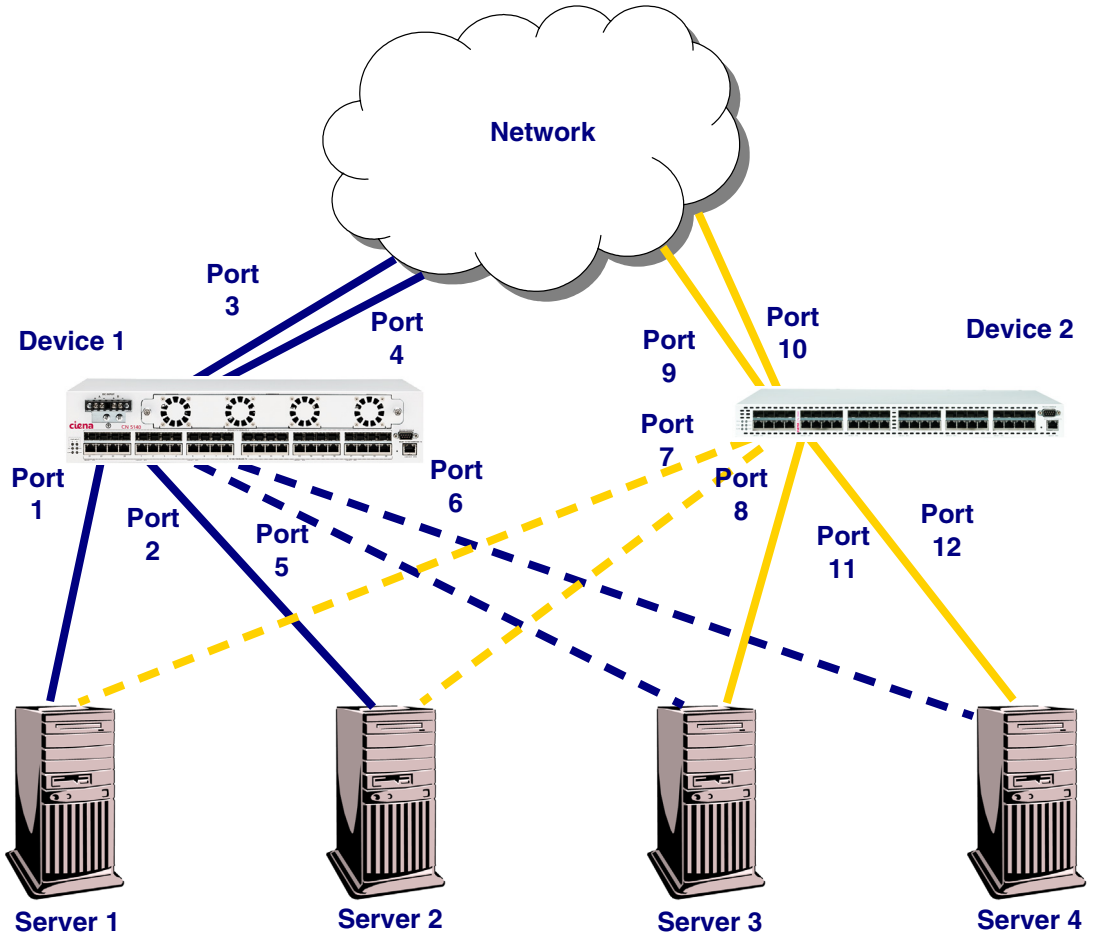

Bi-directional port state mirroring

With bi-directional port state mirroring, ports can mirror their state in both the upstream and downstream directions. This allows the service provider's router port to be automatically disabled if the customer's router port becomes unavailable. Likewise, the converse is true, the customer's router port can be disabled if the service provider's router port becomes unavailable.

This figure shows bi-directional port state mirroring.

Copyright© 2022 Ciena® Corporation

Figure 15-2
Bi-directional port state mirroring

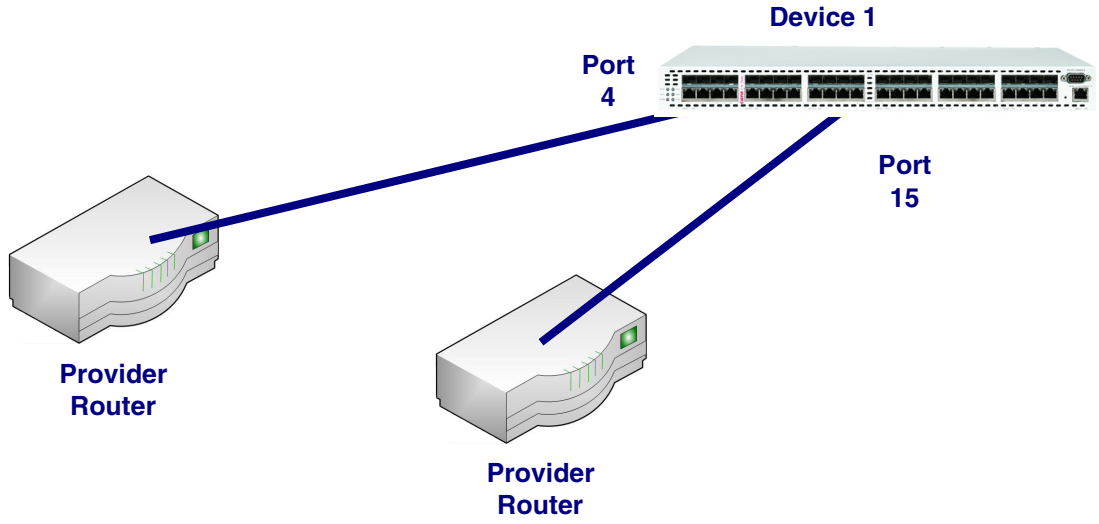

In the Bi-directional port state mirroring figure, assume that port 4 and port 15 have been configured in a bi-directional PSMG. If port 15 cannot establish a link with the customer router, then port 4 becomes disabled as well.

> **CAUTION**
> **Possible Loss of Connectivity**
> This feature must be configured with care in the event that other customers are serviced by the router connected to the device, as they lose connection to the service provider router as well.

Bi-directional PSMGs are configured the same as a standard PSMG; the type must be specified as bi-directional instead of uni-directional.

The procedure for port state mirroring is:

- "Configuring port state mirroring groups" on page 15-11

Limitations

BFD frames generated by the device are not mirrored using egress mirroring if it originates from a hardware accelerated BFD session. This occurs for both traditional and true egress mirroring.

An alternative mirroring solution is to perform ingress mirroring on the port that receives the BFD frames.

Copyright© 2022 Ciena® Corporation

Ingress mirroring and maximum frame size

Several switches have a limitation that pertains to the configuration where the frames switched from ingress port “A” to the egress port “B” are also mirrorred from the ingress port “A” to the egress port “B” using ingress mirroring. Due to hardware limitations of the switching hardware, the maximum frame size on the egress port is disregarded for both the switched frame and for the mirrored frame.

This means that frames ingressing on port “A” that are switched to port “B” and that are normally dropped at the egress port (because they are larger than the maximum frame size on “B”) are actually transmitted when ingress mirroring is configured from “A” to “B”. This limitation applies to these switches:

- 3903
- 3903x
- 3904
- 3905
- 3906

Egress mirroring and maximum frame size

The 3924, 3926 and 3928 switches have a limit on the maximum frame size for packets mirrored with true egress mirroring. A 16-byte header is pushed on to the mirrored frame as it is sent through the internal loopback port. Since the frame size of the mirrored packet is compared against the maximum-frame-size of the egress port and not the mirror-to-port, the mirrored packet is dropped if the original packet exceeds the maximum-frame-size of the egress port minus 16 bytes.

For example, if the maximum frame size of an egress port is 1526 bytes, a packet that is greater than 1510 bytes is not mirrored successfully, regardless of the maximum-frame-size of the mirror-to-port. This limits egress-mirrored frames on these switches to 12,272 bytes. This limitation applies to the 3924, 3926 and 3928 switches.

Copyright© 2022 Ciena® Corporation

Procedure 15-1
Configuring port mirroring

Configure another port to “mirror” the traffic of a specified port to monitor ingress and egress traffic.

Step	Action
1	Enable port mirroring on the mirror port: `port set port <dstPortName> mirror-port on` where <dstPortName> is the port to be configured as the mirror port
2	Configure the source port to be mirrored to mirror ingress and/or egress traffic to the enabled mirror port. `port set port <srcPortName> ingress-mirror <dstportName> egress-mirror <dstPortName>` where <srcPortName> is the port to be mirrored. <dstPortName> is the mirror port.

—end—

Example

This example configures a mirror port to analyze the ingress and egress traffic from port 6 using port 3 as the mirror port.

Enable port mirroring on the mirror port:

```
port set port 3 mirror-port on
```

Configure port 6 to mirror its ingress and egress traffic to port 3:

```
port set port 6 ingress-mirror 3 egress-mirror 3
```


Copyright© 2022 Ciena® Corporation

Procedure 15-2
Configuring encapsulated port mirroring

Configure encapsulated port mirroring to encapsulate mirrored packets with an outer vlan-tag as the mirrored packets egress the mirror-port.

Step	Action
1	Enable port mirroring on the mirror-port: `port set port <dstPortName> mirror-port on` where <dstPortName> is the port to be configured as the mirror port
2	Set the mirroring-encapsulation to vlan-tag and specify the VID and TPID to be used in that tag: `port set port <dstPortName> mirror-encap vlan-tag mirror-encap-vid <vid> mirror-encap-tpid <tpid>` where <dstPortName> is the port to be configured as the mirror port mirror-encap-vid <vid> is the VID to use in the vlan-tag mirror-encap-tpid <tpid> is the TPID to use in the vlan-tag
3	Configure the source port to mirror its ingress and/or egress traffic to the mirror-port: `port set port <srcPortName> ingress-mirror <dstPortName> egress-mirror <dstPortName>` where <srcPortName> is the port to be mirrored. <dstPortName> is the mirror port.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example configures a mirror port to analyze the ingress and egress traffic from port 6 using port 3 as the mirror port and to push a vlan-tag onto the frame as it egresses the mirror-port.

Enable port mirroring on the mirror-port:

```
port set port 3 mirror-port on
```

Set the mirroring-encapsulation to vlan-tag and specify the VID and TPID to be used in that tag:

```
port set port 3 mirror-encap vlan-tag mirror-encap-vid 22 mirror-encap-tpid
8100
```

Configure the source port to mirror its ingress and egress traffic to the mirror-port:

```
port set port 6 ingress-mirror 3 egress-mirror 3
```


Copyright© 2022 Ciena® Corporation

Procedure 15-3
Configuring port state mirroring groups

Configure port state mirroring groups for link redundancy.

Note: Make sure that VLAN/port assignments are configured before attempting this procedure.

Step	Action
1	Create a port state mirror group on the first device: `port state-mirror-group create group <group>` where group <group> — is the port state mirror group.
2	Add source ports and destination ports to <group>: `port state-mirror-group add group <group> dst-port <Port list> src-port <Port list>` where group <group> — is the port state mirror group. dst-port <Port list> — is the destination port. src-port <Port list> — is the source port group.
3	Verify that the ports have been assigned to <group> and that their operational status displays as Link Up: `port state-mirror-group show`
4	Create a port state mirror group on the second device. `port state-mirror-group create group <group2>` where group <group2> — is the port state mirror group on the second device.
5	Add source ports and destination ports to <group2>. `port state-mirror-group add group <group2> dst-port <Port list> src-port <Port list>` where group <group2> — is the port state mirror group on the second device. dst-port <Port list> — is the destination port. src-port <Port list> — is the source port group.

Copyright© 2022 Ciena® Corporation

6 Verify that the ports have been assigned to <group2> and that their operational status displays as Link Up:

```
port state-mirror-group show
```

—end—

Copyright© 2022 Ciena® Corporation

Rapid Spanning Tree Protocol (RSTP) configuration

Spanning Tree Protocol (STP) is a Layer 2 link management protocol defined in IEEE 802.1D. STP provides a loop-free topology in a bridged network and delivers efficient reconfiguration of the loop-free topology in the event that a link fails. Spanning Tree allows devices in an Ethernet network to dynamically create loop-free paths by configuring ports that create loops into a standby (blocking) state.

While link redundancy is desirable for increased fault tolerance, a network topology that has loops (more than one active path to reach end stations) can result in inaccurate forwarding, inaccurate learning, and proliferation of broadcast messages.

Rapid Spanning Tree Protocol (RSTP), formerly defined in IEEE 802.1w and now incorporated in IEEE 802.1D, was developed to achieve faster convergence times. In most circumstances, RSTP can converge the network in less than two seconds. This is critical for networks that carry delay-sensitive traffic such as voice and video. RSTP achieves faster convergence times by selecting not only primary paths to the root bridge, but also backup paths. RSTP is backwards compatible with equipment using the IEEE 802.1D spanning tree algorithm standard.

A primary goal in a Ciena Access Network is to provide failure protection in the 50 ms time frame. This time requirement makes Ciena Carrier Grade Ethernet solutions competitive with SONET failover times as well as meeting the time requirements necessary to protect voice circuits from dropping calls. RSTP reconvergence (change of topology) times can often be below 50 ms.

Copyright© 2022 Ciena® Corporation

- Fault notification - Fault notification is provided by the MEP that detected a connectivity fault for one of these reasons:
  - — expected CCMs were not received
  - — unexpected or invalid CCM were received
  - — the CCM carried a notification of the failure of its associated MEP
- Fault recovery - Fault notifications help network operators correct configuration errors or replace failed components.

This section provides explanations for implementing the Spanning Tree Protocol (STP) and Rapid Spanning Tree Protocol (RSTP) in a Ciena network and also describes Ciena enhancements made to RSTP.

> ***Note:*** This RSTP deployment supports both RSTP and Multiple Spanning Tree Protocol (MSTP). Only one of the two can operate at a time. By default, the mode is RSTP. For details of operating in MSTP mode, see "Multiple Spanning Tree Protocol (MSTP) configuration" on page 17-1.

Ethernet can support many different topologies, such as star, ring, or ladder configurations. While this flexibility makes it ideal for a dynamic network, it can create issues when more than one active path exists to an end station. This is known as a loop.

Copyright© 2022 Ciena® Corporation

For example, in this figure, multiple paths can be taken when forwarding a frame from device A to device Z.

Figure 16-1
Ethernet Topology Example

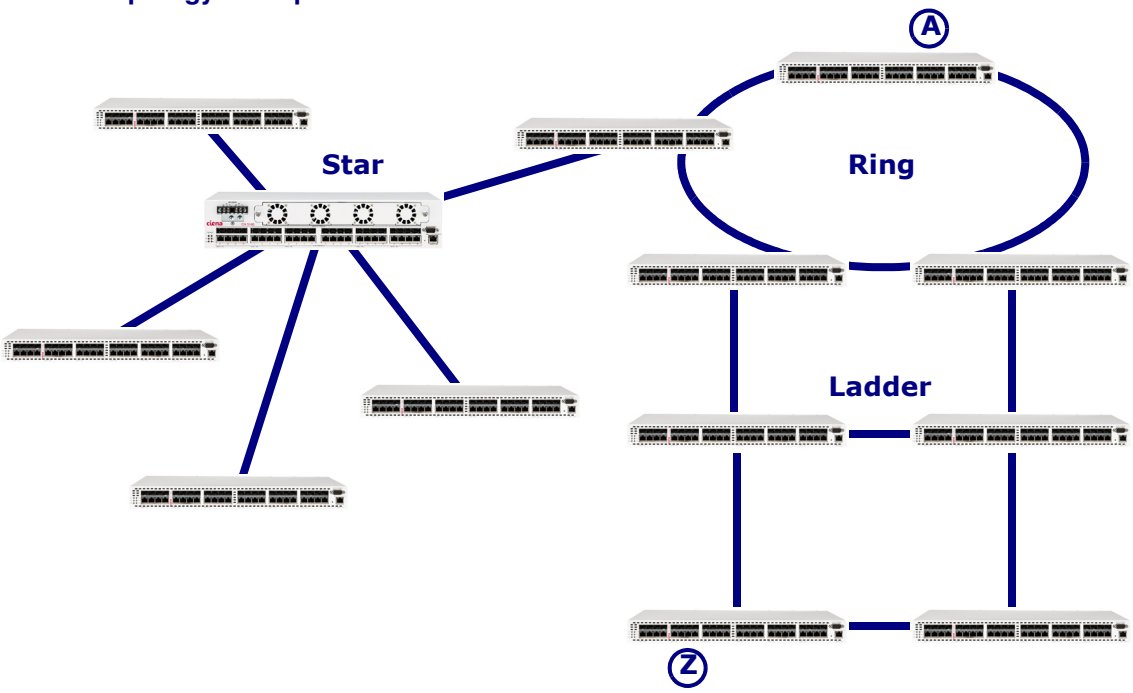

Ethernet traffic forwarding uses Media Access Control (MAC) addresses to forward traffic. When an Ethernet device receives traffic that has a destination MAC address that is unknown to the device, these frames are broadcast to all the devices that it is linked to. This allows the intended destination device to receive the traffic and then reply to the source device. When the first device receives the reply, it then creates a forwarding entry for the destination address. The next time traffic is encountered with that destination address, the device can forward it directly to the correct port without having to broadcast the frame to all egress ports.

This broadcast mechanism is very effective. It can become an issue in network topologies that include physical loops.

Copyright© 2022 Ciena® Corporation

Broadcast storms

A broadcast storm can occur in networks that have loops in their physical topology. A broadcast storm is when large numbers of broadcast packets are unnecessarily propagated through the network, causing unnecessary bandwidth usage and performance degradation.

For example, in the Broadcast Storm figure, if Device A receives a frame with an unknown destination address, it broadcasts the frame out all of its ports in an effort to determine which port the unknown device can be reached on. When Device B receives the frame it also broadcasts the frame to Device C and then C broadcasts it to Device D back to Device A. When Device A receives the frame, it broadcasts the frame again because it still does not have the destination address listed in its table. This can result in the generation of thousands of frames endlessly circling and accumulating in the network until it is completely saturated. STP was created to prevent these broadcast storms.

Figure 16-2
Broadcast Storm

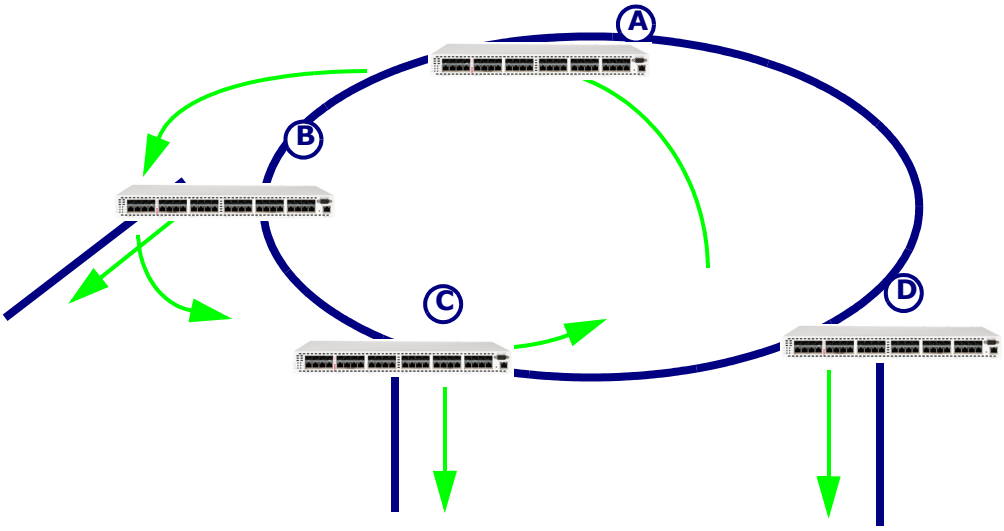

Copyright© 2022 Ciena® Corporation

STP/RSTP goals

For a Layer 2 Ethernet network to function properly, only one active path can exist between any two devices. If multiple paths exist, not only can broadcast storms be generated, but also, duplicate messages might be received by end stations or network devices. These devices may learn duplicate MAC addresses on multiple ports as well. These conditions result in an unstable network. A Spanning Tree Topology is deployed to maintain a stable, loop-free network layout by:

- Electing a root bridge
- Electing the designated bridges and designated ports for each link
- Detecting and preventing loops
- Providing alternate paths

All Spanning Tree topology decisions are made individually by each device. STP/RSTP-enabled devices exchange regular configuration messages known as Bridge Protocol Data Units (BPDUs), with neighboring devices to detect loops, topology changes, elect a root and designated bridges, and recognize the most efficient link path or cost.

Rapid Spanning Tree Protocol implementation

To establish a proper network configuration and topology, the concepts of BPDU, root election, link and path costs, port roles and protocol timers need to be understood. This section reviews these concepts and establishes the recommended proper configuration of the Ciena devices.

Bridge Protocol Data Units

To detect topology changes in the network, RSTP-enabled devices send out bridge protocol data units (BPDUs). BPDUs are layer 2 control frames that carry all information necessary to configure and maintain a loop-free layer 2 RSTP domain. BPDUs are transported using an Ethernet frame with a multicast destination address of 0180C2:000000 and the source MAC address of the port.

BPDUs are sent when a device is powered up, when a topology change is detected, and at regular intervals (every few seconds). Thus, if a device fails, its neighbors realize that they have not received BPDUs and initiate a spanning tree recalculation. When a device receives these RSTP BPDUs, it does not forward them, but uses them to create its own BPDU.

Copyright© 2022 Ciena® Corporation

This table provides a summary of the fields in the BPDU.

Table 16-1
Fields in the BPDU

Field	Description
Port Role	Indicates whether the port is a root port, alternate port, designated port, or a backup port.
Port State	Indicates whether the port is Forwarding, Learning, or Discarding.
RSTP Priority	Specifies a user-configurable value that allows the network operator to select the root bridge.
Root ID	Provides the RSTP priority and base MAC address of the device that the transmitting device considers to be the root.
Root Cost	Provides an additive cost of all the links between the transmitting bridge and the root bridge. A root path cost of 0 indicates this BPDU came directly from the root, each device then adds its link cost towards root to 0 and uses that cost when generating the new BPDU.
Bridge ID	Provides the RSTP priority and base MAC of the transmitting device.
Port Identifier	Provides the port RSTP priority (1-15) and MAC address of the port.
Values for RSTP timers	Provides these values: • message age, which prevents old information from looping endlessly around a failed network. Message age must be less than the max age. • max age, which protects the domain from looping bad information by indicating when a BPDU is considered invalid • hello time, which is the interval at which RSTP BPDUs are sent • forward delay. In RSTP, forward delay is used when a proposal does not get an agree or in half duplex media connections. Once the network is converged, only the root device's timers are used.

Copyright© 2022 Ciena® Corporation

BPDUs are transmitted in one direction, away from the root device. When a device receives a BPDU, it does not forward it, but instead uses the information contained in it to produce its own BPDU. In RSTP, each device sends a BPDU with its current information every Hello Time (by default, 2 seconds). In RSTP, BPDUs serve to converge (reconfigure) the network and to handle topology changes.

Each device in an RSTP domain builds its own configuration of the network by sending BPDUs. The exchange of BPDU transmissions allows each device to determine:

- which device is the root device
- elect the root and designated ports
- the best path to the root device based on the path cost
- alternate paths to the root device
- which ports, if any, must be blocked by STP

Electing the root

BPDUs allow the RSTP devices in the network to elect a root bridge (device). Every Spanning Tree domain has only one root bridge. The root bridge must be the bridge that is closest to the “root” of the physical network. The network administrator must identify this device and configure it as the root bridge.

If no bridge is manually configured as the Root Bridge, on initialization, all bridges consider themselves to be the Root Bridge and send BPDUs out all RSTP-enabled ports. In turn, all RSTP compliant bridges listen for these BPDUs and compare their own bridge ID (RSTP priority and base MAC of the transmitting device) to the ID advertised in the BPDU. Using this process, the device with the lowest bridge ID is selected as the root bridge.

The bridge priority makes up the most significant bytes in the bridge ID. The network operator must manually determine which device becomes the root bridge by configuring the bridge priority. If all bridges are left at the default priority, the device with the lowest MAC address is elected as the root bridge. Link speed, data throughput, and geographic location must also be considered when determining the root bridge.

When selected (automatically or manually), the root bridge sends out configuration messages to the other bridges in the spanning tree. Each bridge then uses this information to calculate the best path to the root bridge and to select which ports are included or blocked in the spanning tree. This process of electing a root bridge and finding out what is connected to every port on the bridge is called convergence.

Copyright© 2022 Ciena® Corporation

Links and path costs

The RSTP negotiation process uses path cost to define root and/or designated devices and ports. To ensure the shortest distance (lowest cost) to the root, a network administrator can configure the switch port's path cost.

The cost of a link is determined or configured by using one of two methods:

- short (IEEE 802.1D, 16-bit format) method
- long (IEEE 802.1t, 32-bit format) method

Ciena devices use the long method by default.

For example, using the short method, a 10 Mbps default cost is 100 and a 100 Mbps default cost is 19.

Table 16-2
Short method: port cost values

Port Speed	Default Cost Value	Range
10 Mbps	100	1 to 65535
100 Mbps	19	1 to 65535
1 Gbps	4	1 to 65535

If the long method is used, a 10 Mbps default cost is 2000000 and a 1 Gbps default cost is 20000.

Table 16-3
Long method: port cost values

Port Speed	Suggested Value	Suggested Range	Available Range
100 Kbps	200000000	20000000 to 200000000	1 to 200000000
1 Mbps	20000000	2000000 to 200000000	1 to 200000000
10 Mbps	2000000	200000 to 20000000	1 to 200000000
100 Mbps	200000	20000 to 2000000	1 to 200000000
1 Gbps	20000	2000 to 200000	1 to 200000000
10 Gbps	2000	200 to 20000	1 to 200000000

Copyright© 2022 Ciena® Corporation

Note: If link aggregation groups are present in the network (a logical link between two devices created by bundling several physical links together), the cost of a link can vary. If a physical link in an aggregation group fails, the amount of traffic that the logical link can carry decreases, increasing its path cost. RSTP calculates path cost based on the maximum logical link capacity, thus does not initiate a recalculation of RSTP topology to minimize the disturbance of the network.

Port states and roles

Once the root bridge has been selected, each bridge port participating in RSTP is placed in one of several different possible roles and states. Port roles and states are determined by the RSTP protocol.

This table lists port states.

Table 16-4
Port states

Port state	Description
Discarding	The port may be RSTP disabled (not transmitting data frames or STP messages), blocking (port is not participating in the active spanning tree topology or data frame transmission), or listening (the port's LAN segment role is being determined by listening and/or sending BPDUs).
Learning	After being in the listening state, a port may transition into the learning state, where it is preparing to enter the forwarding state. The port is not yet included in the active Spanning Tree topology and is still discarding frames.
Forwarding	The Spanning tree is in the converged state and the port is actively forwarding frames.

Copyright© 2022 Ciena® Corporation

This table provides a summary of RSTP port Roles/States.

Table 16-5
RSTP port role and state summary

RSTP port role	Port state	Description
Root Port	Forwarding	Represents the best path (lowest cost) to the Root.
Alternate Port	Discarding (Blocking)	Secondary port (Root Port being primary) that provides an alternate path to the root bridge if the root port fails.
Designated Port	Forwarding	Port on the designated bridge which incurs the lowest path cost and is used to forward traffic away from the root to the destination bridge.
Backup Port	Discarding (Blocking)	Secondary port (Designated Port being primary) that provides an alternate path to a LAN segment if the Designated Port fails.
Loopback Port	Discarding (Blocked)	Bridge port that is in a loopback mode and does not take part in active topology computation.

A root port is selected on each RSTP device (with the exception of the root device). The root port is the port through which the designated device forwards frames toward the root device. Each designated device has one unique root port. In this device, traffic is forwarded from the root port to one or more of the designated ports and vice versa, following the device forwarding configuration. All the other ports are inactive, either disabled or blocking. The RSTP port priority and port path cost setting determine which ports are set in the forwarding state and which ports are set in the blocking state. The RSTP port priority value represents the location of a port in the network topology and how well it is able to forward traffic.

After the root device is selected, RSTP determines a designated bridge and a designated port for each destination device. The designated bridge is the only device that is allowed to forward traffic from the root bridge to a particular destination. The designated port is the port used to forward traffic away from the root to a destination device.

A port with the root or a designated port role is included in the active RSTP topology. A port with an alternate or backup port role is excluded from the active RSTP topology. It is only activated and becomes part of the RSTP topology in case of a failure of the root or designated port.

Copyright© 2022 Ciena® Corporation

Protocol timers

RSTP uses several timers. These timers can be changed but are generally set at their default values.

This table lists the timers that can be configured on Ciena devices.

Table 16-6
Timers

Timer	Description
Forward-delay	Indicates in seconds the time the bridge spends in the learning state. The default value is 15 with a possible range of 4 to 30.
Max-age	Sets the maximum length of time in seconds a port saves its valid lifetime configuration message. The default value is 20 with a possible range of 6 to 40.
Hello-time	Indicates the time in seconds between BPDUs transmitted between bridges. The default value is 2 with a range of 1 to 10.

It is recommended to keep these values at their default factory settings. If a change of topology occurs during the transmission of a configuration message (BPDUs) in a large network involving many bridges, the reconvergence might never configure the correct topology. If these outdated configuration messages are not updated quickly, a device might cancel or reverse the reconvergence in progress, thus initiating another reconvergence.

To counter this possible scenario, each BPDU is flagged with a Message Age parameter, incremented with each passage through a device. The Max-age validates the lifetime of the BPDU and flags it as null or not valid when the Max-age value is reached. An improper setting of these aging parameters might not detect outdated BPDUs and prevent the network from properly converging.

Copyright© 2022 Ciena® Corporation

Converging the network with RSTP

The main goal of RSTP is to remove network loops and converge the network as quickly as possible. The network is said to be converged when all RSTP devices have determined their roles/port states, and traffic can traverse the network in any direction. This means that the root device and designated devices have been determined and that loops are artificially disabled by putting ports into a blocking state. A port that is in a blocking state does not transmit any traffic, and drops all receiving traffic except incoming BPDUs. The root bridge, along with the designated bridges, becomes the Spanning Tree.

In the past, STP took, on average, 30 to 40 seconds to converge the network. RSTP has shaved that time down to less than 2 seconds, on average. With Ciena RSTP implementation, convergence times have been reduced to less than 200 ms, on average, and often are closer to 50 ms.

Topology changes

Once the Spanning Tree has been established, RSTP monitors the network for failures and topology changes. Events that can cause the topology to change include device failure, link failure, new links, new devices, or new device configurations. As these events occur, RSTP automatically adapts the tree to the new topology. After the spanning tree adjusts in response to a topology change, the network is said to be re-converged.

IEEE 802.1D STP was designed to support enterprise networks where a recovery from a failure in several minutes was acceptable. To improve the time needed to reconfigure the spanning tree, RSTP introduced several mechanisms to improve the re-convergence time. These mechanisms allow the spanning tree to perform rapid transitions from the current port state.

IEEE 802.1w RSTP introduces the concept of edge ports. Edge ports are ports on spanning tree devices that connect to an edge device. An edge device is one that does not participate in the spanning tree (doesn't acknowledge or send BPDUs). Edge ports are those ports that have not received BPDUs.

Figure 16-3
Edge Port

Copyright© 2022 Ciena® Corporation

An edge port transitions to the forwarding state as soon as it receives a BPDU from the edge device, without first listening and learning. Edge ports also do not send a topology change notification when they transition to forwarding.

RSTP also uses alternate ports to increase the speed of re-convergence. Alternate ports are “standby” root ports. In the event that the link connecting a root port fails, the device has an Alternate Port that becomes the Root Port and transitions to the forwarding state, often in less than 50 ms. This is represented in the Alternate Port Switch-over figure, where if the link for the Root Port of Device A fails, an Alternate Port takes over as the Root Port.

Figure 16-4
Alternate Port Switch-over

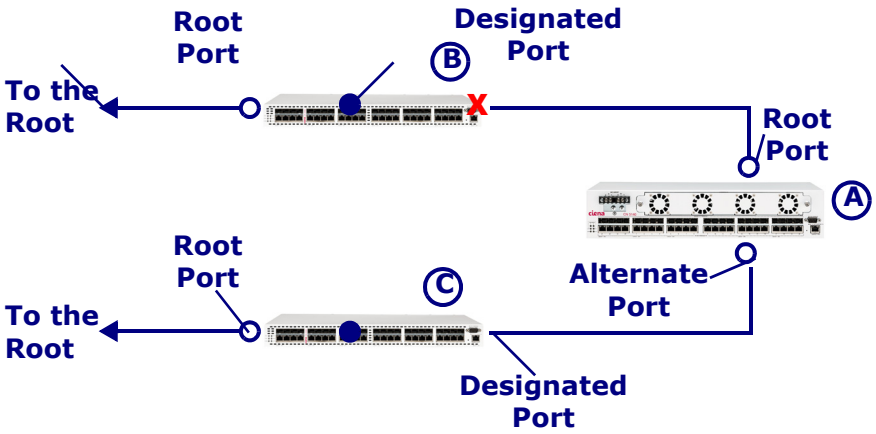

Copyright© 2022 Ciena® Corporation

Another mechanism for improving RSTP re-convergence speed is the rapid transition of designated ports using a Proposal-Sync-Agreement mechanism. In the Proposal-Sync-Agree Mechanism figure, the designated port on bridge A is directly connected to the Root Port on Bridge C. When Device A detects a topology change, the designated port enters the discarding or learning state and sets a proposal bit on BPDUs that it sends out.

The direction of the root device is determined and Device C puts its non-edge designated ports into a blocking state. Device C is then considered to be “synchronized” and sends a BPDU toward the root with the “agreement” bit set. On receiving the agreement from the downstream bridge, the designated port on Device A then enters the forwarding state. This whole process occurs with the exchange of a few frames.

Figure 16-5
Proposal-Sync-Agree Mechanism

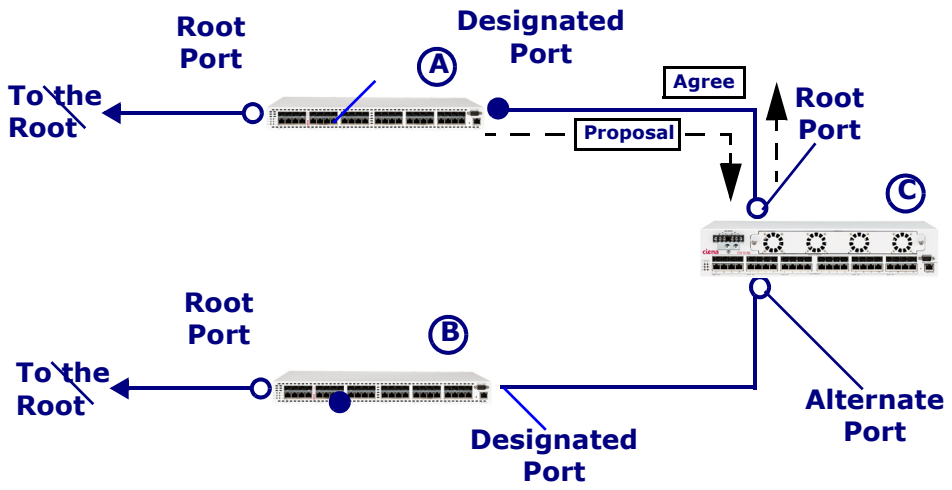

Copyright© 2022 Ciena® Corporation

Spanning tree domains

RSTP is a global protocol. It runs across a bridged Ethernet network. As the access network grows, it may be desirable to limit the reach of RSTP to certain areas of the network. An example of a network topology that can benefit from configuring RSTP is shown in the Spanning tree domains figure. The provider network is carried on the main ring and services are then extended to edge rings connected to the core provider ring.

Figure 16-6
Spanning tree domains

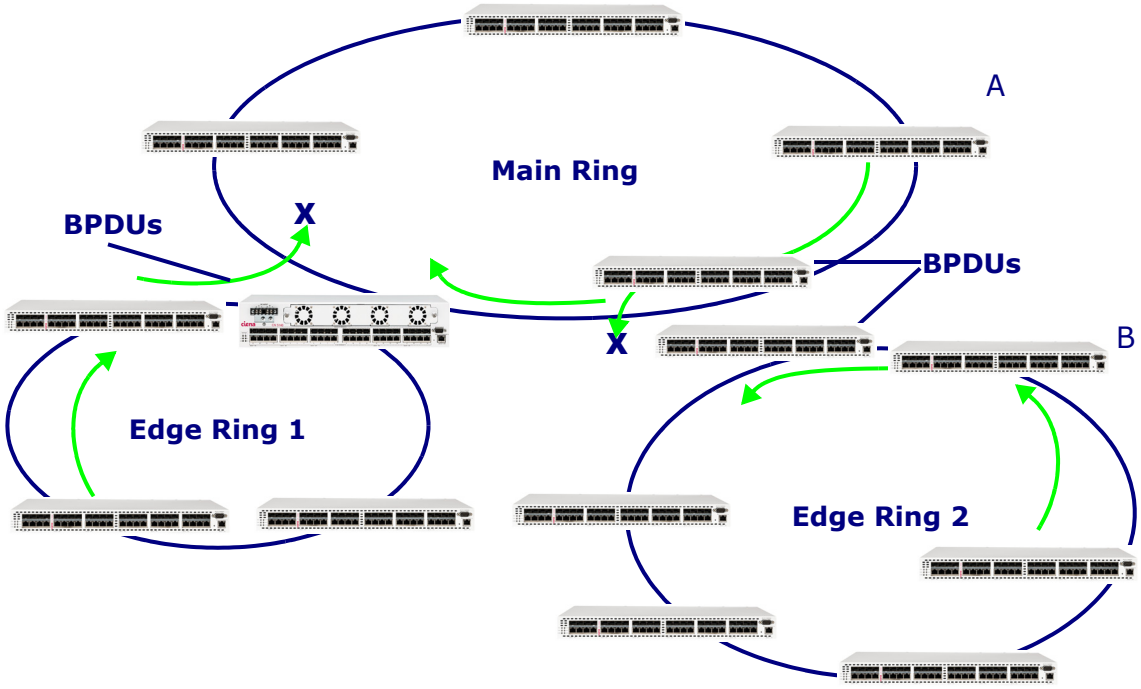

If a topology change occurs on an edge ring, the reconvergence event can impact the main provider ring, and cause traffic problems across the entire network. To minimize the impact that topology changes have on the network, Spanning Tree Domains can be defined.

Traffic in and between these domains flows without any restriction. As seen in Figure 16-6, RSTP BPDUs are only forwarded between devices in the same Spanning Tree Domain. Edge Ring 1 is one Spanning Tree domain and Edge Ring 2 and the Main Ring also form their own Spanning Trees. This prevents a link failure on Edge Ring 2 from causing traffic problems on the other rings.

Copyright© 2022 Ciena® Corporation

Edge rings can also be configured using Multi-Device Spanning Tree Domains, which allow an edge ring to connect to two devices on the main ring, increasing reliability and redundancy. Spanning Tree Domains and Multi-Device Spanning Tree Domains are fully compliant with the IEEE 802.1D and 802.1w standards as the format of the BPDUs exchanged is not modified. See "Multiple Spanning Tree Protocol (MSTP) configuration" on page 17-1 for more information.

By default, each device has a single Spanning Tree domain. Additional Spanning Tree domains are then created and ports are moved from the default domain to the new domain. The system software supports up to 32 RSTP domains. A port can only be a member of one Spanning Tree domain, but any number of ports can be added to an RSTP domain, up to the maximum number of ports on the device. Spanning Tree domains are only defined at the local level, this means that other devices in the network are not aware that their neighbors are partitioned. Any port assigned to a domain is able to move to any other domain. The default domain name is 0, and domain names up to 31 characters are supported.

> ***Note:*** Configuring multiple RSTP domains with multiple bridges interconnecting the domains causes the domains to merge. In Figure 16-6, connecting device A to device B initiates the merging of the main domain ring and the edge domain ring 2. The system software has no way to prevent this misconfiguration.

Loop detection and prevention

Another RSTP enhancement is the detection and prevention of loops caused by non-bridge network devices, for example, hubs, connected to a Ciena device. After sending a BPDU, if a Ciena device detects its own BPDU coming back, it means that a physical loop exists on that part of the network. The Ciena device automatically changes the RSTP state to BLOCK, and sends an SNMP (Simple Network Management Protocol) trap to alert the operator of the loop condition. RSTP then periodically tries to re-enable RSTP on the port to detect an eventual suppression of that loopback.

Copyright© 2022 Ciena® Corporation

RSTP standards compliance

While RSTP is designed to provide rapid convergence, its speed depends on the network topology and the location of the failure in the network. The Ciena solution is designed to support the delivery of time-critical services such as voice and video. To maintain acceptable service levels during topology changes, reconvergence must occur in 200 milliseconds or less. This is much faster than RSTP was designed to support.

Ciena RSTP enhancements are compatible and interoperable with IEEE 802.1D STP and IEEE 802.1w RSTP, and operate in a mixed vendor network deployment without any modification to the BPDUs. The overall reconvergence performance of the network is significantly enhanced when Ciena devices are exclusively implemented.

You can enable Ciena RSTP enhancements when running in RSTP or RSTP domain mode. In RSTP domain mode, standards compliance is configured separately for each domain. There are three separate compliance levels to select from as shown in this table.

Table 16-7
RSTP compliance levels

Compliance Settings	Standard
ciena-enhanced (default)	802.1D-2004 with proprietary enhancements
8021d	Strict 802.1D-2004
8021w	Strict 802.1w

Compliance Modes 8021D and ciena-enhanced can support up to 41 node rings.

Procedures for RSTP are:

- “Changing RSTP standards compliance” on page 16-18
- “CLI configuration example” on page 16-19

Copyright© 2022 Ciena® Corporation

Procedure 16-1
Changing RSTP standards compliance

You can:

- change the compliance mode
- change the compliance setting in RSTP domain mode

Step	Action

To change the compliance mode

1 Change the compliance mode:

```
rstp set compliance <8021w|8021D|ciena-enhanced>
```

where

compliance <8021w|8021D| ciena-enhanced>	is the RSTP standards compliance level.

To change the compliance setting in RSTP domain mode

2 Change the compliance setting in RSTP domain mode:

```
rstp set rstpd <rstp> compliance <8021w|8021D|ciena-enhanced>
```

where

rstpd <RstpDomain>	is the RSTP domain name.
compliance <8021w|8021D| ciena-enhanced>	is the RSTP standards compliance level.
priority <NUMBER: 0-15>	is the port RSTP priority.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 16-2
CLI configuration example

In this example, the devices A, B, and C are connected to ensure redundancy and prevent loops.

This figure shows the sample RSTP configuration.

Figure 16-7
Sample RSTP configuration

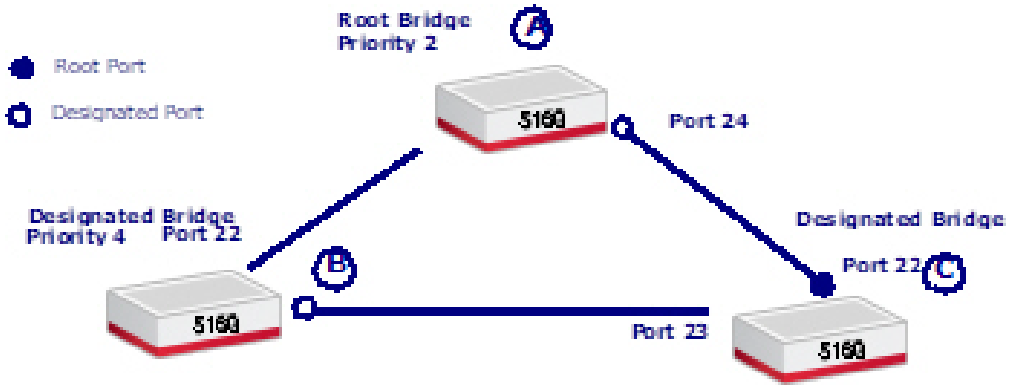

This table lists the configuration for the devices.

Table 16-8
Devices configuration

Device	Bridge type	Bridge priority	Port	Port priority	Port size	Port type	Port cost
A 5160	Root	2	23 24	6 8	Gig Gig	Designated Designated	20000 20000
B 5160	Designated	4	22 23	4 6	Gig Gig	Root Designated	20000 20000
C 5160	Designated	6	22 23	4 6	Gig Gig	Root Alternate	20000 20000

Copyright© 2022 Ciena® Corporation

To configure the RSTP example (for each device):

Step	Action
1	Set the mode to RSTP domain mode. `rstp set mode rstp-domain`
2	Set the device priority. `rstp set [priority <NUMBER:0-15>]` where priority <NUMBER:0-15> is the bridge RSTP priority.
3	Create an RSTP domain. `rstp create rstpd <String{15]> [domain-id <NUMBER: 1-255>]` where rstpd <String[15]> is the RSTP domain name. domain-id <NUMBER: 1-255> is the RSTP domain identifier.
4	Add ports to the domain. `rstp assign rstpd <rstp> port <Port list>` where rstpd <rstp> is the RSTP domain name. port <Port list> is the port list.
5	Set port priority. `rstp set port <port> [priority <NUMBER:0-15>]` where port <port> is the port list. priority <NUMBER:0-15> is the port RSTP priority.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example creates the RSTP configuration shown in Figure 16-7.

Configuration of Device A:

```
rstp set mode rstp-domain
rstp set priority 2
rstp create rstpd RSTP-Domain
rstp assign rstpd RSTP-Domain port 23,24
rstp set port 23 priority 4
rstp set port 24 priority 6
```

Configuration of Device B:

```
rstp set mode rstp-domain
rstp set priority 4
rstp create rstpd RSTP-Domain
rstp assign rstpd RSTP-Domain port 22,23
rstp set port 22 priority 4
rstp set port 23 priority 6
```

Configuration of Device C:

```
rstp set mode rstp-domain
rstp set priority 6
rstp create rstpd RSTP-Domain
rstp assign rstpd RSTP-Domain port 22,23
rstp set port 22 priority 4
> rstp set port 23 priority 6
```


Copyright© 2022 Ciena® Corporation

Copyright© 2022 Ciena® Corporation

Multiple Spanning Tree Protocol (MSTP) configuration

This section explains the implementation of the Multiple Spanning Tree Protocol (MSTP), including these topics:

- "Overview" on page 17-1
- "Configuration" on page 17-9

Overview

Multiple Spanning Tree Protocol (MSTP - IEEE 802.1Q-2005) is a standards-based version of creating multiple spanning trees where each VLAN has its own Multiple Spanning Tree Instance (MSTI). Each instance has an independent spanning tree topology to provide multiple forwarding paths for data traffic, enable load balancing and reduce the number of spanning tree instances required to support a large number of VLANs. MSTP can be implemented in place of Rapid Spanning Tree Protocol (RSTP).

Like RSTP, defined in IEEE 802.1D, MSTP provides a loop-free topology in bridged networks and delivers efficient reconfiguration (convergence) of the loop-free topology in the event that a link fails. MSTP inherits its rapid transition mechanism from RSTP, defined in IEEE 802.1W, to achieve fast convergence times.

Multiple Spanning Tree Protocol (MSTP) provides:

- Scalability on large networks. Both STP and RSTP 802.1W allow the total BPDU traversing of maximum 24 bridges from the root bridge. RSTP 802.1D/D4 (2003) allows the BPDU traversing of maximum 40 bridges. In an MST region, MSTP limits BPDU traversing of maximum of 30 bridges, with no limit on the number of regions. In case of ring topology, the number of bridges a BPDU can traverse is limited to the maximum number of bridges configured on the ring, taking the worst case of link failure into consideration. In a meshed topology, the maximum number of bridges that can be included in the topology is much more than the limit.

Copyright© 2022 Ciena® Corporation

- Limited Topology changes. MSTP limits the topology change events as in the region as possible. This makes a lot of sense, especially for large networks.
- Redundant links and load balancing. MSTP supports the assignment of a number of VLANs to a single spanning tree instance, and multiple instances are supported in each spanning tree region. These multiple Spanning Tree instances enable L2 networks to have redundant physical links because VLANs can be assigned to different Virtual bridges (spanning tree instances); this also enables load balancing.

MST regions and spanning trees

An MST Region is a set of physically connected LANs and MST bridges that support the same MSTP configuration. A region can have one bridge or multiple bridges. Each bridge must be capable of processing RSTP BPDUs. There is no limit to the number of MST regions in a network, and each region supports up to 8 MSTIs. One of these MSTIs is called an Internal Spanning Tree (IST), which connects all MSTP bridges in an MST Region. An IST is a sub tree in the Common Spanning Tree (CST) that spans the network and interconnects MST Regions. The collection of ISTs and the CST make up the Common and Internal Spanning Tree (CIST).

Bridges

The bridge with the best priority among all STP, RSTP, and MSTP enabled network bridges is called the CIST Root. The boundary bridge between regions becomes the CIST Regional Root. If there are multiple boundary bridges, the one with the best cost to the CIST Root becomes the CIST Regional Root. If all the bridges in a network are in one MST region, the CIST Root is also the CIST Regional Root.

BPDUs

Like its predecessors, STP and RSTP, MSTP-enabled devices send out bridge protocol data units (BPDUs) to detect topology changes in the network. BPDUs are sent when a device is powered up, when a topology change is detected, and at regular intervals (every few seconds). Thus, if a device fails, its neighbors realize that they have not received BPDUs and initiate a spanning tree recalculation. When a device receives these BPDUs, it does not forward them, but uses them to create its own BPDUs.

Copyright© 2022 Ciena® Corporation

MSTP BPDUs carry the information listed in this table.

Table 17-1
MSTP BPDUs

Field	Description
Protocol Identifier	Identifies the protocol as STP, RSTP, or MSTP.
Protocol Version Identifier	Identifies the protocol version as STP (0), RSTP (2), or MSTP (3).
BPDU Type	Specifies configuration or topology change MSTP/RSTP BPDU.
CIST Flags	Identifies the message as a topology change, proposal, port role, learning, forwarding, agreement, or topology change acknowledgment.
CIST Root Identifier	Identifies the bridge priority and base MAC address of the device that the transmitting device considers to be the CIST Root.
CIST External Root Path Cost	Provides an additive cost of all the links between the transmitting bridge and the CIST Root bridge. A root path cost of 0 indicates this BPDU came from the same region as the CIST Root, each device then adds its link cost towards root to 0 and uses that cost when generating the new BPDU.
CIST Regional Root Identifier	Provides the bridge priority and base MAC address of the device that the transmitting device considers to be the CIST Regional Root.
CIST Port Identifier	Provides the port priority and base MAC address of the port transmitting the BPDU.
Forward Delay	Forward delay is the time taken before a root or designated port starts forwarding after being selected and determined not to satisfy the rapid transition condition. The operational value is set to the value distributed by the root bridge via BPDU and the configured value is not set until this bridge becomes the elected root bridge. The default value is 15.
Message Age	Specifies the current age of the BPDU.
Max Age	Specifies the BPDU life span.
Hello Time	Specifies the time between two consecutive periodic BPDUs.

Copyright© 2022 Ciena® Corporation

Table 17-1
MSTP BPDUs

Field	Description
Version 1 Length	Indicates whether or not Version 1 (IEEE Standard 802.1G) information is present. For MSTP, the value is 0 to indicate that no Version 1 information is present.
Version 3 Length	Specifies the number of bytes that are to follow in the packet.
MST Configuration Identifier	Provides the “fingerprint” of the MSTP configuration for the bridge, including: • Configuration Id Format Selector - The format of the MST Configuration Identifier, which is 0, the value defined for MSTP in the 802.1Q-2005 standard. • Configuration Name - Configurable character string that uniquely identifies the region. The default value is the device’s MAC address. • Revision Level - Configurable numeric value that provides an additional parameter to further differentiate MST regions. • Configuration Digest - A 16-octet signature that is calculated with the HMAC_MD5 algorithm from the configured VLAN ID and MSTID mapping table as specified in the standard.
CIST Internal Root Path Cost	Specifies how close the device is to the CIST Root bridge.

Copyright© 2022 Ciena® Corporation

Table 17-1
MSTP BPDUs

Field	Description
CIST Bridge Identifier	Provides the base MAC and bridge priority of the CIST Bridge.
CIST Remaining Hops	Specifies the number of hops remaining before reaching the maximum hop value.
MSTI Configuration Messages	Provides information used by MSTP to calculate vectors for the specified multiple spanning tree instances (MSTIs), including: • MSTI Flags - Identifies the message as a topology change, proposal, port role, learning, forwarding, agreement, or master. • MSTI Regional Root Identifier - base MAC and bridge priority of the MSTI Regional Root. • MSTI Internal Root Path Cost - an additive cost of all the links between the transmitting bridge and the MSTI Regional Root bridge. An MSTI Regional Root path cost of 0 indicates this BPDU came directly from the MSTI Regional Root, each device then adds its link cost towards MSTI Regional Root to 0 and uses that cost when generating the new BPDU. • MSTI Bridge Priority - bridge priority value. • MSTI Port Priority - port priority value. • MSTI Remaining Hops - hops remaining before the number exceeds the maximum hop count.

Port roles

Based on the MSTP calculation for load balancing, ports in the CIST and each MSTI are assigned port roles. The port roles are similar to the port roles in STP and RSTP except for the CIST Regional Root Port role which is only for MSTP.

Copyright© 2022 Ciena® Corporation

This table lists port roles.

Table 17-2
MSTP Port Role and State Summary

MSTP Port Role	Port State	Description
CIST Root Port	Forwarding	Represents the best path (lowest cost) to the Root bridge of the CIST/MSTI of the bridge.
Alternate Port	Discarding (Blocking)	Secondary port (Root Port being primary) that provides an alternate path to the root bridge of he CIST/MSTI if the root port fails.
Designated Port	Forwarding	Port that provides the best path to the root bridge of the CIST/MSTI for the link.
Backup Port	Discarding (Blocking)	Port that provides a backup path to a LAN segment if the Designated Port of the CIST or MSTI fails.
CIST Regional Root Port	Forwarding	For MSTI only, this port provides connectivity from the Region to a CIST Root that lies outside the Region.

Note: This implementation does not support the learning state as defined by the IEEE. When a port is behaving as defined by the IEEE learning state, it is in the blocking state in this implementation.

Boundary port

A boundary port connects an MST region to:

- A different MST region
- A single STP box
- A single RSTP box

At the boundary port of the region, MSTP configuration changes so that the control of the STP state is per port rather than per port per VLAN, since the MSTI for the VLAN cannot map to the boundary port. For this reason, MSTP calculated roles do not apply to its boundary ports. Instead, the boundary port state is forced to be the same as the IST port state.

Copyright© 2022 Ciena® Corporation

Network convergence

When a MSTP bridge initializes, it claims itself as both the CIST Root and CIST Regional Root via BPDU. When superior CIST Root information is received, it relinquishes its claim as the CIST Root. When superior IST information is received, it then relinquishes its claim as the CIST Regional Root. Finally, the MST region is converged when all bridges agree on the same CIST Regional Root and all port roles are synchronized.

MSTP limits the scope of topology change. If any MSTI topology change is detected by the non-boundary ports, it only triggers a MAC flush to the MSTI in the region. A topology change detected by the boundary port triggers MAC flushes more widely.

When the network is converged, each MST region independently allocates any frames to any given VLAN in a stable network. The CIST Root, Designated, and CIST Regional Root ports forward data frames, and Alternate, Backup and Disabled ports do not.

Backward compatibility

As designed by IEEE, MSTP is fully backward compatible with RSTP and STP. MSTP operates with switches that support STP, RSTP, and traditional 802.1q through the selection of a designated MSTP bridge to function as the representative of that entire region to the CST. This enables MSTP regions composed of many switches and bridges to appear as a single switch to other Ethernet switches and bridges in a network. BPDUs for the STP instances are exchanged only between switches in an MSTP region. Since this information is not propagated outside the MSTP region, switch CPU cycles are reduced. Also, keeping this information in the region prevents network flapping and frequent network reconvergence caused by link and port outages.

The MSTP uses version 3 BPDU which attaches the CIST and MSTI information at the end of legacy STP and RSTP BPDUs. The CST takes care of the connectivity for all MST regions and all the RSTP and STP bridges in the network. The MST region appears as a virtual bridge cluster to adjacent STP or RSTP bridges (as well as to other MST regions.)

Copyright© 2022 Ciena® Corporation

When an MST enabled bridge port receives an STP BPDU (version 0 BPDU) or an RSTP BPDU (version 2 BPDU), it knows that it is at the boundary of the region. Based on the version of the BPDU received, an MST enabled bridge sends the corresponding version of BPDU to communicate with its counterpart. For interoperability with STP, only STP BPDUs are sent out, and for RSTP there is no protocol migration required since RSTP can decode CIST information as RSTP BPDU. When its neighbor bridge switches its force-version to MSTP, it does not force this link to MSTP automatically. A per port protocol migration check forces the port to MSTP for a short period of time.

> ***Note:*** MSTP backward compatibility supports RSTP mode and STP mode. It does not support Ciena proprietary RSTP port based domain functionality.

Copyright© 2022 Ciena® Corporation

Configuration

Assume a network with two regions, each with three bridges, separated by an RSTP bridge.

Figure 17-1
Sample Network Segment

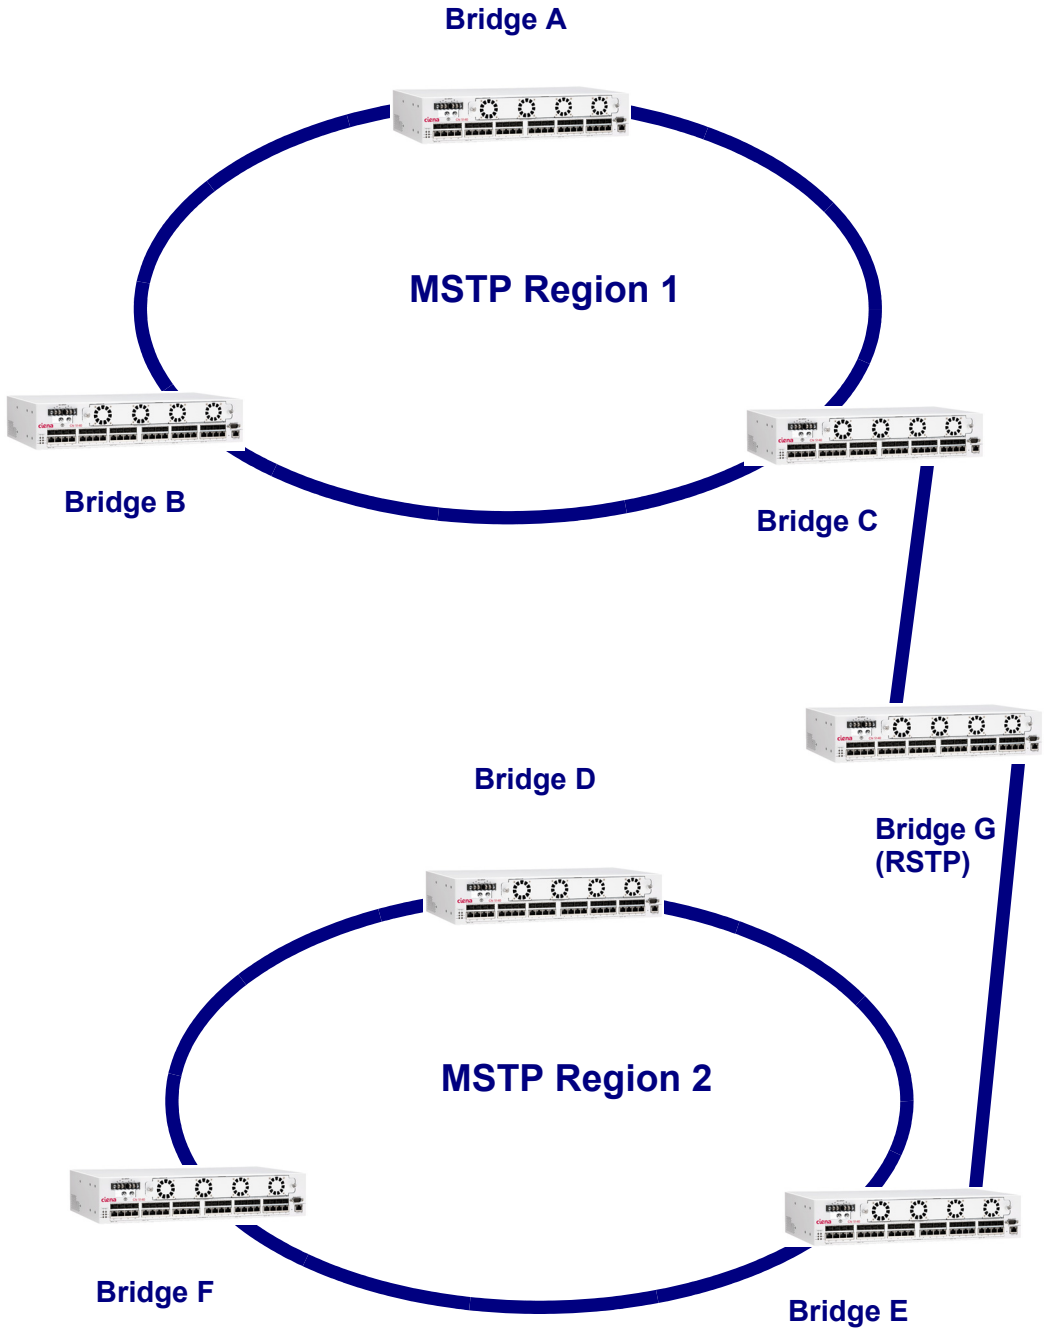

Copyright© 2022 Ciena® Corporation

The general steps for configuring MSTP in a network are as follows:

1 Configure one bridge in the network as the CIST root bridge.
2 Configure the CIST port settings.
3 For each region, configure the unique MST Identifiers.
4 For each MSTP bridge, configure the mode to MSTP.
5 Configure the MSTIs and MSTI port settings.
6 Configure the bridge settings.
7 Configure port settings.
8 Enable MSTP on all bridges.

Configuring the CIST root bridge

For all bridges in the network, pick a bridge as the root bridge of the CIST. Assign the lowest number of priority to the root bridge.The bridge-priority assigned to the CIST for the bridge is used in the priority component of the CIST Bridge Identifier of BPDUs transmitted from this bridge. The supported range is 0-15. The value entered is multiplied by 4096 to conform with the standard as shown in this table.

Table 17-3
CIST Root Bridge Priority Mapping

Bridge Priority Entered Value	Maps to Standard Value
0	0
1	4096
2	8192
3	12288
4	16384
5	20480
6	24576
7	28672
8	32768
9	36864
10	40960
11	45056
12	49152

Copyright© 2022 Ciena® Corporation

Table 17-3
CIST Root Bridge Priority Mapping (continued)

Bridge Priority Entered Value	Maps to Standard Value
13	53248
14	57344
15	61440

Procedures for MSTP are:

- "Configuring MSTP CIST attributes" on page 17-12
- "Configuring CIST ports" on page 17-13
- "Configuring the MST Identifier" on page 17-15
- "Changing mode" on page 17-17
- "Configuring MSTIs" on page 17-19
- "Changing the MSTI bridge priority" on page 17-23
- "Deleting an MSTI" on page 17-25
- "Configuring MSTI ports" on page 17-26
- "Configuring bridge global settings" on page 17-28
- "Configuring ports" on page 17-30
- "Enabling and disabling MSTP on specified ports" on page 17-32
- "Enabling MSTP and saving configuration" on page 17-33
- "Disabling and enabling MSTP on a bridge" on page 17-34
- "Disabling and enabling MSTP on an MSTI" on page 17-35
- "Displaying MSTP information" on page 17-36
- "Clearing MSTP statistics" on page 17-39

Copyright© 2022 Ciena® Corporation

Procedure 17-1
Configuring MSTP CIST attributes

Configure MSTP CIST attributes.

Step	Action
1	Configure MSTP CIST attributes: `mstp set cist <cist> bridge-priority <NUMBER: 0-15>` where cist <cist> — is the CIST identifier. bridge-priority <NUMBER: 0-15> — is the CIST bridge-priority.
2	Verify the priority: `mstp show cist`

—end—

Example

Assume for the network segment shown in Figure 17-1 on page 17-9 that Bridge A is the CIST Root bridge.

This example sets the CIST bridge priority for Bridge A.

```
mstp set cist bridge-priority 1
```

This example verifies the priority.

```
> mstp show cist

+--- MSTP CIST Info -------------------------------------------------------+
| Parameter                          | Value                                |
+------------------------------------+--------------------------------------
| CIST Bridge Id                     |                                      |
|   Priority                         | 1  (0x1000)                          |
|   Mac Address                      | 00-02-A1-30-7C-40                    |
...
+------------------------------------+--------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 17-2
Configuring CIST ports

Configure CIST ports.

Port priority is the priority component of the CIST Port Identifier field of BPDUs transmitted from the specified port(s). The supported range is 0-15. The value entered is multiplied by 16 to conform with the standard as shown in this table.

Table 17-4
CIST port priority mapping

Bridge priority entered values	Maps to standard values
0	0
1	16
2	32
3	48
4	64
5	80
6	96
7	112
8	128
9	144
10	160
11	176
12	192
13	208
14	224
15	240

Copyright© 2022 Ciena® Corporation

Step	Action
1	Configure CIST attributes for a port: `mstp set cist-port <cist-port> port-priority <NUMBER: 0-15> int-path-cost <NUMBER: 1-200000000> dynamic-int-path-cost <on|off> {port <Port list>}`

where

cist-port <cist-port>	is the CIST identifier.
port-priority <NUMBER: 0-15>	is the CIST port priority.
int-path-cost <NUMBER: 1-200000000>	is the CIST internal port path cost. Sets the CIST internal root path cost for the specified port(s). Allowable path cost values are 1-200000000. This sets the CIST Internal Root Path Cost field in BPDUs transmitted from the specified port(s). This value is not used in the BPDU if dynamic internal path cost is enabled for the CIST on the specified port(s). The default is 20000.
dynamic-int-path-cost <on|off>	specifies whether dynamic internal port cost is on or off. Enables or disables dynamic internal port path cost calculations for the CIST for the specified port(s). When enabled, the CIST Internal Root Path Cost field of BPDUs from the specified port(s) use a value internally calculated based on link rate (and possibly other factors). When disabled, the configured int-path-cost value for the CIST for the specified port(s) is used.
port <Port list>	is the port list.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 17-3
Configuring the MST Identifier

The MST Configuration Name and Revision Level must match for all the nodes in the same region. The default MST Configuration Name is the bridge MAC. All bridges in one region, must be configured to same value. Otherwise, each bridge acts as if its in different region.

Step	Action
1	Configure the MST Identifier: `mstp set [mst-cfg-id-name <String[32]>] [mst-cfg-id-rev-level <NUMBER: 0-65535>]` where mst-cfg-id-name <String[32]> — is the MSTP configuration ID name. The configuration name is used for calculating the MST Configuration Identifier for the bridge. mst-cfg-id-rev-level <NUMBER: 0-65535> — is the MSTP configuration ID revision level. The numeric value is transmitted in MST BPDUs in the Revision Level of the MST Configuration Identifier field. Since this is limited to 2 octets, the range of this value is 0-65535. The default value is '0'.
2	Verify the MST identifier: `mstp show mst-cfg-id`

—end—

Example

These examples configure the network segment shown in Figure 17-1 on page 17-9:

This example configures the MST Identifier for Bridges A, B, and C.

```
mstp set mst-cfg-id-name "MST Region 1" mst-cfg-id-rev-
level 1
```


Copyright© 2022 Ciena® Corporation

This example verifies the MST Identifier.

```
mstp show mst-cfg-id

+--- MSTP MST Configuration Id Info -----------------------------------+
| Parameter                          | Value                            |
+------------------------------------+----------------------------------+
| Configuration Name                 |                     MST Region 1 |
| Revision Level                     |                                1 |
| Configuration Digest               | 6EDEDB2C31B7A50DCEBFB6FFE47D5766 |
+------------------------------------+----------------------------------+
```

This example sets the MST Identifier for Bridges D, E, and F.

```
mstp set mst-cfg-id-name "MST Region 2" mst-cfg-id-rev-
level 2
```

This example verifies the MST Identifier.

```
mstp show mst-cfg-id

+--- MSTP MST Configuration Id Info -----------------------------------+
| Parameter                          | Value                            |
+------------------------------------+----------------------------------+
| Configuration Name                 |                     MST Region 2 |
| Revision Level                     |                                2 |
| Configuration Digest               | 6F696C6D9DD9032F8FEDA4DD25AC021B |
+------------------------------------+----------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 17-4
Changing mode

Selecting the mode specifies the version of Spanning Tree for the bridge. This MSTP deployment supports both legacy RSTP and MSTP. Only one of the two can operate at a time. Once the device is powered up in an STP mode and operating in that mode, you can still configure all the settings for the alternate mode, and then re-boot to immediately start operation in the alternate mode.

The initialization to operate in RSTP or MSTP mode occurs during the boot-up sequence. By default, the mode is RSTP.

For details of operating in RSTP mode, see "Rapid Spanning Tree Protocol (RSTP) configuration" on page 16-1.

To configure the network segment shown in Figure 17-1 on page 17-9, you need to change the mode for Bridges A, B, C, D, E, and F.

Step	Action
To change the mode to MSTP	
1	Determine the current mode: `mstp show bridge`
2	If the mode is RSTP, change the mode to MSTP: `mstp set mode mstp`
3	Run the MSTP configuration commands for provisioning the device in MSTP mode:
4	Save the configuration: `configuration save`
5	Reboot the device: `reboot`
6	After the reboot, verify that the mode has changed: `mstp show bridge`
To change the mode to RSTP:	
7	Determine the current mode: `rstp show bridge`
8	If the mode is MSTP, change the mode to RSTP: `mstp set mode rstp`
9	Run the RSTP configuration commands for provisioning the device in RSTP mode:

Copyright© 2022 Ciena® Corporation

10 Save the configuration:

```
configuration save
```

11 Reboot the device:

```
reboot
```

12 After the reboot, verify that the mode has changed:

```
rstp show bridge
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 17-5
Configuring MSTIs

Configure an MSTI if you want to create an MSTI for a specific VLAN. By default, all VLANs are mapped to the CIST.

Note these limitations:

- The system software currently enforces MSTI identifier values 1-64. The standard supports values 1-4094. Devices cannot participate in MST regions configured with an MSTI identifier value greater than 64.
- There is a limit of 16 MSTIs. To provide maximum flexibility, any MSTI index in the allowable range of 1-64 may be created, but once more than 16 instances have been created, an error message is displayed when additional MSTI creations are attempted.
- A VLAN cannot be part of both Ethernet Ring Protection Switching (G.8032) and MSTP.
- If a VLAN is already added to a virtual ring, it can not be assigned to an MSTI. If that VLAN needs to be assigned to an MSTI, then remove the VLAN from the virtual ring, and then add it to the MSTI.
- If a VLAN is already assigned to MSTI, it can not be added to a virtual ring. If that VLAN needs to be added to a virtual ring, reassign it to the CIST and then add it to the virtual ring.
- While enabling G.8032 and MSTP on the same device, care must be taken to assign ports only to the VLANs that fall under respective MSTP or G.8032 enabled regions. If a MSTP enabled VLAN includes a port from the G.8032 region, then loops can occur. Similarly, if a G.8032 VLAN has a port in an MSTP region, then loops can occur.

Step	Action
1	Create the VLAN: `vlan create vlan <vlan> [name <String[31]>]` where vlan <vlan> — is the VLAN ID. name <String[31]> — is the name of the VLAN.

Copyright© 2022 Ciena® Corporation

2 Add the specific ports to the VLAN:

```
vlan add vlan <vlan> port <Port list>
```

where

vlan <vlan>	is the VLAN list to add ports.
port <Port list>	is the list of ports to add to VLAN.

3 Create the MSTIs:

```
mstp create msti <NUMBER: 1-64> [bridge-priority <NUMBER:
0-15>]
```

where

msti <NUMBER: 1-64>	is the MSTI. A maximum of 16 MSTIs is allowed.
bridge-priority <NUMBER: 0-15>	is the MSTP bridge priority.

4 Display the MSTI

```
mstp show msti <NUMBER: 1-64>
```

where

msti <NUMBER: 1-64>	is the MSTI.

5 Map the VLAN to the MSTI.

```
mstp set mst-cfg-id-vlan-mapping <<mst-cfg-id-vlan-
mapping> {msti <NUMBER: 1-64>}
```

where

mst-cfg-id-vlan-mapping <mst-cfg-id-vlan-mapping>	is the MSTP configuration ID VLAN mapping.
msti <NUMBER: 1-64>	is the MSTI that the VLAN is to be mapped to.

6 Display MST configuration mapping information:

```
mstp show mst-cfg-id <mst-cfg-id> vlan-mapping-table xst-
mapping-table
```

where

mst-cfg-id <mst-cfg-id>	is the MST ID.
vlan-mapping-table	is the mst-id VLAN-mapping.
xst-mapping-table	is the mst-id xst-mapping-table.

—end—

Copyright© 2022 Ciena® Corporation

Example

Create the VLAN.

```
vlan create vlan 1001-1003
```

Add the specific ports to the VLAN.

```
vlan add vlan 1001 port 1-5
vlan add vlan 1002 port 6-10
vlan add vlan 1003 port 11-15
```

Create the MSTIs.

```
mstp create msti 1
mstp create msti 2
mstp create msti 3
```

Display the MSTI.

```
> mstp show msti 1
+--- MSTP MSTI 1 Info -------------------------------------------------------+
| Parameter                          | Value                                  |
+------------------------------------+----------------------------------------+
| MST ID                             | 1                                      |
| MSTI Enable                        | enabled                                |
|                                    |                                        |
| MSTI Bridge Id                     |                                        |
|   Priority                         | 8  (0x8001)                            |
|   Mac Address                      | 00-03-18-BA-66-80                      |
|                                    |                                        |
| MSTI Bridge Priority Vector        |                                        |
|  Regional Root Id                  |                                        |
|   Priority                         | 8  (0x8001)                            |
|   Mac Address                      | 00-03-18-BA-66-80                      |
|  Internal Path Cost                | 0                                      |
|  Designated Bridge Id              |                                        |
|   Priority                         | 8  (0x8001)                            |
|   Mac Address                      | 00-03-18-BA-66-80                      |
|  Designated Root Id                | (0x0000)                               |
|   Priority                         | 0                                      |
|   Port                             | 0                                      |
|                                    |                                        |
| MSTI Bridge Times                  |                                        |
|  Max Hops                          | 20                                     |
|                                    |                                        |
| MSTI Root Priority Vector          |                                        |
|  Regional Root Id                  |                                        |
|   Priority                         | 8  (0x8001)                            |
|   Mac Address                      | 00-03-18-BA-66-80                      |
|  Internal Path Cost                | 0                                      |
|  Designated Bridge Id              |                                        |
|   Priority                         | 8  (0x8001)                            |
|   Mac Address                      | 00-03-18-BA-66-80                      |
|  Designated Root Id                | (0x0000)                               |
|   Priority                         | 0                                      |
```


Copyright© 2022 Ciena® Corporation

```
|   Port                              | 0                                        |
|                                     |                                          |
| MSTI Root Port ID                   | (0x0000)                                 |
|   Priority                          | 0                                        |
|   Port                              | 0                                        |
|                                     |                                          |
| MSTI Root Times                     |                                          |
|  Remaining Hops                     | 20                                       |
+-------------------------------------+------------------------------------------+
```

Map the VLAN to the MSTI.

```
mstp set mst-cfg-id-vlan-mapping 1001 msti 1
mstp set mst-cfg-id-vlan-mapping 1002 msti 2
mstp set mst-cfg-id-vlan-mapping 1003 msti 3
```

Display the XST mapping table.

```
> mstp show mst-cfg-id xst-mapping-table

+--- MSTP MST Configuration Id Xst Mapping -------------------------------------+
| Xst      | Vlans                                                                 |
+----------+-----------------------------------------------------------------------+
| CIST     | 1-4094                                                                |
| MSTI  1  | (none)                                                                |
+----------+-----------------------------------------------------------------------+
```

Display the VLAN mapping table.

```
> mstp show mst-cfg-id vlan-mapping-table

+--- MSTP MST Configuration Id Vlan Table -------------------------------------+
|Vlan|Xst| Vlan|Xst| Vlan|Xst| Vlan|Xst| Vlan|Xst| Vlan|Xst| Vlan|Xst| Vlan|Xst|
+----+---+-----+---+-----+---+-----+---+-----+---+-----+---+-----+---+-----+---+
|   1|CST|  513|CST| 1025|CST| 1537|CST| 2049|CST| 2561|CST| 3073|CST| 3585|CST|
|   2|CST|  514|CST| 1026|CST| 1538|CST| 2050|CST| 2562|CST| 3074|CST| 3586|CST|
|   3|CST|  515|CST| 1027|CST| 1539|CST| 2051|CST| 2563|CST| 3075|CST| 3587|CST|
|   4|CST|  516|CST| 1028|CST| 1540|CST| 2052|CST| 2564|CST| 3076|CST| 3588|CST|
|   5|CST|  517|CST| 1029|CST| 1541|CST| 2053|CST| 2565|CST| 3077|CST| 3589|CST|
|   6|CST|  518|CST| 1030|CST| 1542|CST| 2054|CST| 2566|CST| 3078|CST| 3590|CST|
|   7|CST|  519|CST| 1031|CST| 1543|CST| 2055|CST| 2567|CST| 3079|CST| 3591|CST|
...
+----+---+-----+---+-----+---+-----+---+-----+---+-----+---+-----+---+-----+---+
```


Copyright© 2022 Ciena® Corporation

Procedure 17-6
Changing the MSTI bridge priority

The bridge-priority assigned to the specified MSTI for the bridge is sent in the MSTI Bridge Priority field of the associated MSTI Configuration Message for the specified MSTI in BPDUs transmitted from this bridge. The specified MSTI must be that of a created MSTI (in the range of 1-64).

The supported range for MSTI Bridge Priority is 0-15. The value entered is multiplied by 4096 to conform with the standard as shown in this table.

Table 17-5
MSTI bridge priority mapping

Bridge priority entered value	Maps to standard value
0	0
1	4096
2	8192
3	12288
4	16384
5	20480
6	24576
7	28672
8	32768
9	36864
10	40960
11	45056
12	49152
13	53248
14	57344
15	61440

Copyright© 2022 Ciena® Corporation

Step	Action
1	Change the MSTI bridge priority: `mstp set msti <msti> bridge-priority <NUMBER: 0-15>` where msti <msti> — is the MSTI identifier. bridge-priority <NUMBER: 0-15> — is the MSTI bridge priority. **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 17-7
Deleting an MSTI

Delete an MSTI when it is no longer required.

Step	Action
1	Remove the VLAN mapping from the MSTI. `mstp set mst-cfg-id-vlan-mapping <mst-cfg-id-vlan-mapping> cist` where mst-cfg-id-vlan-mapping <mst-cfg-id-vlan-mapping> — is the MSTP configuration ID VLAN mapping. cist — maps the VLAN to the CIST
2	Delete the MSTI. `mstp delete msti <NUMBER: 1-64>` where msti <NUMBER: 1-64> — is the MSTI to be deleted.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 17-8
Configuring MSTI ports

Configure MSTI ports to set priority and path cost settings for MSTI ports.

The MSTI port priority mapping table lists MSTI port priority mapping. The supported range is 0 to 15. The value entered is multiplied by 16 to conform with the standard.

Table 17-6
MSTI port priority mapping

Bridge priority entered value	Maps to standard value
0	0
1	16
2	32
3	48
4	64
5	80
6	96
7	112
8	128
9	144
10	160
11	176
12	192
13	208
14	224
15	240

Copyright© 2022 Ciena® Corporation

Step	Action
1	Configure the MSTI ports:

```
mstp set msti-port <msti-port> port-priority <NUMBER: 0-15> int-path-cost <NUMBER: 1-200000000> dynamic-int-path-cost <on|off> {port <Port list>}
```

where

msti-port <msti-port>	is the MSTI identifier.
port-priority <NUMBER: 0-15>	is the MSTI port priority. Sets the MSTI port-priority for the specified MSTI for the specified port(s). Allowable priority values are 0-15. Allowable msti-port values are any valid (created) MSTI (in the range of 1-64). This sets the MSTI Port Priority field used in the associated MSTI Configuration Message contained in MST BPDUs transmitted from the specified port(s).The supported range is 0-15. The value entered is multiplied by 16 to conform with the standard as shown in Table 17-6 on page 17-26.
int-path-cost <NUMBER: 1-200000000>	is the MSTI internal port path cost. Sets the MSTI internal root path cost for the specified MSTI for the specified port(s). Allowable path cost values are 1-200000000. Allowable MSTI ID values are any valid (created) MSTI (in the range of 1-64). This sets the MSTI Internal Root Path Cost field in the associated MSTI Configuration Message in BPDUs transmitted from the specified port(s). This value is not used in the BPDU if dynamic internal path cost is enabled for the specified MSTI on the specified port(s). The default is 20000.
dynamic-int-path-cost <on|off>	determines whether dynamic internal port cost is on or off. Enables or disables dynamic internal port path cost calculations for the specified MSTI for the specified port(s). When enabled, the MSTI Internal Root Path Cost field of the associated MSTI Configuration Message in MST BPDUs transmitted from the specified port(s) use a value internally calculated based on link rate (and possibly other factors). When disabled, the configured int-path-cost value for the specified MSTI for the specified port(s) is used.
port <Port list>	is the port list.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 17-9
Configuring bridge global settings

Configure bridge global settings.

Step	Action
1	Configure bridge global settings: `mstp set forward-delay <NUMBER: 4-30> hello-time <NUMBER: 1-10> loopback-blocking <on|off> max-age <NUMBER: 6-40> max-hops <NUMBER: 6-40> mode <mstp | rstp> path-cost-def <stp8021d1998|stp8021t2001> tx-hold-count <NUMBER: 1-10> mst-cfg-id-name <String[32]> mst-cfg-id-rev-level <NUMBER: 0-65535> pvst-notification <on | off>` where forward-delay <NUMBER: 4-30> — set bridge MSTP forward delay timer Forward delay is the time taken before a root or designated port starts forwarding after being selected and determined not to satisfy the rapid transition condition. The operational value is set to the value distributed by the root bridge via BPDU and the configured value is not set until this bridge becomes the elected root bridge. The default value is '15'. hello-time <NUMBER: 1-10> — set bridge MSTP hello time Hello time is the time between two consecutive periodic BPDUs.The operational value is set to the value distributed by the root bridge via BPDU and the configured value is not set until this bridge becomes the elected root bridge. The minimum value of '1' is not recommended because the periodic BPDU sending takes up the 1 per second transmit hold count allowance. When rapid convergence takes place, only one BPDU is allowed to send in a second for a port delaying the explicit handshake and rapid topology change. The default value is '2'. loopback-blocking <onloff> — set bridge MSTP loopback port blocking Loopback port blocking prevents data loops by detecting their own BPDUs being looped back to them on the same port. The default value is 'off', which is shown as 'disabled' with the mstp show bridge command. max-age <NUMBER: 6-40> — set bridge MSTP loopback port blocking Maximum age limits the BPDU life span to limit the size of the spanning tree. The default value is '20'

Copyright© 2022 Ciena® Corporation

where

max-hops <NUMBER: 6-40>	set bridge MSTP maximum hops Maximum hops limits the number of hops for BPDUs. The default is '20'.
mode <mstp | rstp>	set bridge MSTP mode
path-cost-def <stp8021d1998|stp8021t2001>	set bridge MSTP path-cost default Path cost default controls the type of internal and external path costs for the bridge. MSTP supports default path costs defined in IEEE802.1D-1998 and IEEE802.1t-2001. The default value is 'stp8021t2001'.
tx-hold-count <NUMBER: 1-10>	set bridge MSTP TX hold count Transmit hold count limits the number of BPDUs a bridge can send per second. The default value is '6'.
mst-cfg-id-name <String[32]>	MSTP config id name
mst-cfg-id-rev-level <NUMBER: 0-65535>	MSTP config ID revision level
pvst-notification <on | off>	turn on or off PVST notifications

—end—

Copyright© 2022 Ciena® Corporation

Procedure 17-10
Configuring ports

Configure ports.

Step	Action
1	Configure ports: `mstp set port <port> dynamic-ext-path-cost <on|off> edge-port <on|off> migration-check ext-path-cost <NUMBER: 1-200000000> point-point-mac <auto|force-false|force-true> auto-edge <on|off> restricted-role <on|off> restricted-tcn <on|off>` where port <port> — is the port list. dynamic-ext-path-cost <on|off> — turn dynamic external port cost on/off Turns on or off dynamic calculation of the external path cost. When set to 'on', the external port path cost is calculated dynamically based on the port speed. When set to 'off', the external port path cost is the value configured in the 'ext-path-cost' parameter. edge-port <on|off> — turn port edge port state on/off Setting the specified port to be consider an edge port allows for rapid port state transition to the forwarding state. For 10/100 ports, the default value for edge-port is 'on'. For other ports, the default is 'off'. Always configure GIG link edge ports explicitly as edge ports. Rapid transition fails otherwise. Another alternative is to disable MSTP on the GIG link edge ports, though this is not the preferred method. migration-check — turn port migration check on/off Triggers a migration check by the MSTP Port Protocol Migration State Machine. This is a momentary configuration parameter, setting it to 'on' initiates the migration check and setting it to 'off' has no effect. ext-path-cost <NUMBER: 1-200000000> — set port MSTP external path cost Sets the external port path cost for the specified port, which takes effect only if 'dynamic-ext-path-cost' is set to 'off'. The default value is '20000'. Rather than configure the port path cost, it is recommended that you turn on dynamic path cost to have MSTP select the best link(s) for forwarding.

Copyright© 2022 Ciena® Corporation

where

point-point-mac <auto|force-false|force-true>	set point-to-point MAC mode Sets how to determine the link for the specified port: auto (automatic mechanisms determine the link), force-true (connected to a point-to-point link) or force-false (not connected to a point-to-point link).
auto-edge <on|off>	turn auto edge on/off Sets MSTP to automatically identify the edge ports.
restricted-role <on|off>	turn restricted role on/off When set to on, this attribute restricts the port from becoming a root port, even if it has the best priority. The default is 'off'.
restricted-tcn <on|off>	turn restricted tcn on/off When set to on, this attribute restricts the port from forwarding topology change notices.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 17-11
Enabling and disabling MSTP on specified ports

You can:

- enable MSTP on specified ports
- disable MSTP on specified ports

Step	Action

To enable MSTP on specified ports

1 Enable MSTP on specified ports:

```
mstp enable port <port list>
```

where

port <port list> is the port list.

To disable MSTP on specified ports

2 Enable MSTP on specified ports:

```
mstp disable port <port list>
```

where

port <port list> is the port list.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 17-12
Enabling MSTP and saving configuration

Enable MSTP on all bridges and save the configuration after completing the MSTP configuration.

Step	Action
1	Globally enable MSTP: `mstp enable`
2	Save the configuration: `configuration save`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 17-13
Disabling and enabling MSTP on a bridge

Disable MSTP on a bridge when required for troubleshooting. Enable MSTP on a bridge when the problem is resolved.

> ***Note:*** It is best never to disable MSTP on a bridge in the network as data loops might exist, rendering part or whole of the network useless.

	Step	Action
To disable MSTP		
	1	Disable MSTP: `mstp disable`
To enable MSTP		
	2	Enable MSTP: `mstp enable`

—end—

Copyright© 2022 Ciena® Corporation

Procedure 17-14
Disabling and enabling MSTP on an MSTI

Disable MSTP on an MSTI when required for troubleshooting. Enable MSTP on an MSTI when the problem is resolved.

When an MSTI is disabled, its identifier is not be transmitted in MSTI Configuration Messages in the MST BPDUs transmitted from the bridge.

> ***Note:*** Use caution when disabling MSTP as data loops can exist in the network topology.

Step	Action

To disable MSTP on a specified MSTI

1 Disable MSTP on a specified MSTI:

```
mstp disable msti <NUMBER: 1-64>
```

where

msti <NUMBER: 1-64> is the MSTI identifier.

To enable MSTP on a specified MSTI

2 Enable MSTP on a specified MSTI:

```
mstp enable msti <NUMBER: 1-64>
```

where

msti <NUMBER: 1-64> is the MSTI identifier.

—end—

Copyright© 2022 Ciena® Corporation

Procedure 17-15
Displaying MSTP information

You can display:

- MSTP information
- MSTP port information
- MTSI port information
- MST identifier information

Step	Action

To display MSTP information

1 Display MSTP information:

```
mstp show bridge cist statistics
```

where

bridge	displays bridge information.
cist	displays CIST information.
statistics	displays MSTP statistics information. MSTP statistics track the number of received and transmitted topology change notification (TCN), configuration (CFG), RSTP, and MSTP BPDUs.

To display MTSP port information

2 Display MSTP port information:

```
mstp show port <port> statistics state-machine
```

where

port <port>	is the port list.
statistics	displays port statistics.
state-machine	displays port state machine. MSTP port and CIST state machine information includes variables, timers, and states.

To display MTSI port information

3 Display MSTI port information:

```
mstp show msti <msti>
```

where

msti <msti>	is the port list.

Copyright© 2022 Ciena® Corporation

To display MST identifier information

4 Display MST identifier information:

```
mstp show mst-cfg-id <mst-cfg-id> vlan-mapping-table xst-mapping-table
```

where

mst-cfg-id <mst-cfg-id>	is the MST identifier.
vlan-mapping-table	displays mst-id VLAN-mapping.
xst-mapping-table	displays mst-id xst-mapping-table.

—end—

Example

This example shows sample output for the mstp show statistics command.

```
> mstp show statistics

+--- MSTP Statistics -------------------------------+
| Port   | Port Receive Statistics                   |
| Name   | rxTcnBpdu| rxCfgBpdu| rxRstBpdu| rxMstBpdu|
+--------+----------+----------+----------+----------+
|1       |         0|         0|         0| 909098983|
|2       |         0|         0|         0|         0|
|3       |         0|         0|         0|         0|
|4       |         0|         0|         0|         0|
...
+--------+----------+----------+----------+----------+
| Port   | Port Transmit Statistics                  |
| Name   | txTcnBpdu| txCfgBpdu| txRstBpdu| txMstBpdu|
+--------+----------+----------+----------+----------+
|1       |         0|         0|         0|    120889|
|2       |         0|         0|         0|         0|
|3       |         0|         0|         0|         0|
|4       |         0|         0|         0|         0|
...
+--------+----------+----------+----------+----------+
```


Copyright© 2022 Ciena® Corporation

This example shows sample output for the show port statistics command.

```
> mstp show port 1 statistics

+--- MSTP Port 1 Statistics ---------------------------------------------+
| Statistic                          | Value                               |
+------------------------------------+-------------------------------------+
| Received TCN BPDUs                 | 0                                   |
| Received CFG BPDUs                 | 0                                   |
| Received RST BPDUs                 | 3                                   |
| Received MST BPDUs                 | 3                                   |
|                                    |                                     |
| Transmitted TCN BPDUs              | 0                                   |
| Transmitted CFG BPDUs              | 0                                   |
| Transmitted RST BPDUs              | 0                                   |
| Transmitted MST BPDUs              | 36627                               |
+------------------------------------+-------------------------------------+
```

This example shows sample output for the mstp show port state machine command.

```
> mstp show port 1 state-machine

+--- MSTP Port 1 Info ---------------------------------------------------+
| Parameter                          | Value                               |
+------------------------------------+-------------------------------------+
| Name                               | 1                                   |
| Port Status                        | up                                  |
| Port Uptime                        | 0000D19:52:40                       |
|                                    |                                     |
| MSTP Enable                        | enabled                             |
...
+------------------------------------+-------------------------------------

+--- MSTP Port 1 CIST Info ---------+-------------------------------------+
| MSTP State                         | forwarding                          |
| Role                               | designated                          |
...
+------------------------------------+-------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 17-16
Clearing MSTP statistics

Clear MSTP statistics.

Step	Action
1	Clear statistics for all or specified ports: `mstp clear port <port> [statistics]` where port <port> is the port list. statistics clears statistics. **—end—**

Copyright© 2022 Ciena® Corporation

Copyright© 2022 Ciena® Corporation

Ethernet Ring Protection switching configuration

This section explains the configuration of Ethernet Ring Protection (ERP) switching.

Configuring ERP switching comprises configuring major rings and optional sub-rings. R-APS channels cannot be shared between a sub-ring and a major ring: their messages must be restricted to their own rings.

A major ring is a complete ring of G.8032 nodes, where each node is connected to two other nodes. The R-APS channel messages of all virtual rings are communicated to all nodes in the ring.

Sub-rings connect to two nodes of the major ring. The R-APS channel of the sub-ring is terminated at the connecting endpoint nodes by specifying the termination points for the sub-ring.

Each link and port can be a member of only one ring, either a sub-ring or a major ring.

Every major ring and every sub-ring must have its own ring protection link configured to provide a channel block.

R-APS messages in the G.8032 standard use a destination MAC address in the 01:09:A7:00:00:XX range to communicate between ring nodes. In Ciena packet networking devices, the last byte of this address identifies the logical ring and is equal to the configured ring ID of the logical ring. The VID specified in the R-APS messages are specific to the R-APS VID configured for the virtual ring. So both the destination MAC address and the VID in a R-APS message identify a unique virtual ring on the device.

When interworking with other vendors in the same G.8032 logical ring, it is recommended that the ring ID on the Ciena packet networking device is set to the same value that matches the last byte of the destination MAC address of the R-APS messages that the other vendor is using.

Copyright© 2022 Ciena® Corporation

For the number of ports, logical rings, and virtual rings by platform, refer to the table in the “G.8032 logical and virtual rings” topic in the “System capabilities” section of *39XX/51XX Planning, Engineering and Ordering*.

MC-LAG interworking with G.8032

If MC-LAG is configured with G.8032, then flush propagation must be enabled for MC-LAG (RG nodes) sub-ring/major nodes.

```
>ring-protection virtual-ring set ring flush-propagate <off|on>
```

Procedures

This figure shows the flow of ERP switching configuration procedures.

Figure 18-1
ERP switching configuration

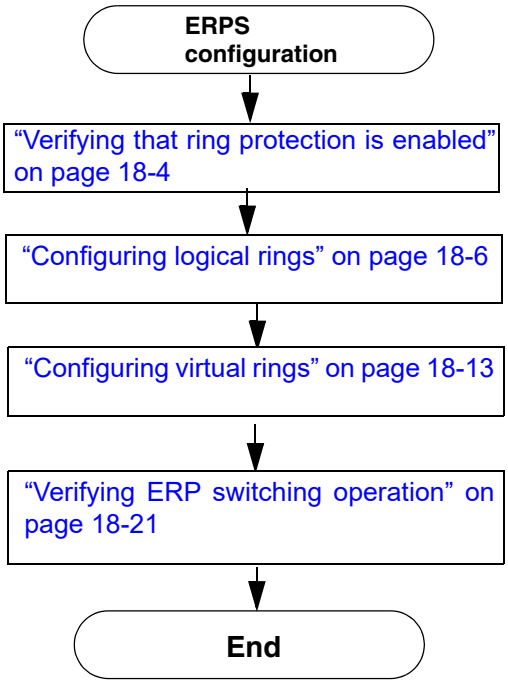

Procedures to configure ERP switching are

- “Verifying that ring protection is enabled” on page 18-4
- “Configuring logical rings” on page 18-6
- “Configuring virtual rings” on page 18-13
- “Verifying ERP switching operation” on page 18-21

Copyright© 2022 Ciena® Corporation

Optional procedures are

- “Displaying information for logical rings” on page 18-22
- “Displaying virtual ring information” on page 18-24
- “Displaying ring protection” on page 18-27
- “Displaying ERP switching in the configuration file” on page 18-28
- “Checking configuration for a ring not forwarding traffic” on page 18-29

Additionally, this section provides a “Configuration example” on page 18-30.

Copyright© 2022 Ciena® Corporation

Procedure 18-1
Verifying that ring protection is enabled

By default, ring protection is enabled. Make sure that ring protection is enabled before configuring ERP switching.

Step	Action
1	Display ring protection: `ring-protection show`
2	If the value of the State parameter is Disabled, enable ring protection: `ring-protection enable`

—end—

Example

This example shows sample output for the ring-protection show command. In this example, ring protection is enabled.

```
> ring-protection show

+--------------- RING-PROTECTION INFO ----------------+
| Parameter                    | Value                |
+------------------------------+----------------------+
| State                        | Enabled              |
| Node ID                      | 00:02:5A:01:B0:AA    |
| Ether Type                   | 0x8100               |
| Number Of Logical Rings      | 1                    |
+------------------------------+----------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 18-2
Disabling ring protection

There may be occasions during troubleshooting that require you to globally disable ring protection.

Step	Action
1	Disable ring protection: `ring-protection disable` **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 18-3
Configuring logical rings

Configure a logical ring to define the port on the switch that is connected to two other switches in a ring topology to transport the traffic ERPS is intended to protect. The switch is connected to two other switches to form a closed loop.

A logical ring is defined by the attributes listed in this table.

Table 18-1
Attributes of logical rings

Attribute	Description
Name	Unique alpha-numeric string for the logical ring.
Ring Identifier	Numeric identifier 1-255. SAOS supports "stacked" G.8032 logical rings with a common logical ring ID. This means: • SAOS allows the user to create more than one logical ring with the same logical ring ID. All virtual rings supported by a given node (no matter what ring ports are being used) are uniquely identified by the <logicalRingID, raps-VID> pair. • All virtual rings sharing a common logical ring port can have the same logical ring identifier (that is, <ringID>), and thus are uniquely identified by the <ringID, raps-VID> pair. For example, if the <ringID> is the same (for example, has a value of 1) across all logical rings, then the <raps-VID> needs to be unique.
Guard-Time	Prevents nodes that are transmitting recovery messages from acting on outdated R-APS messages. Range is 10-2000 milliseconds. Default is 500ms.
Wait to Restore (WTR)	Time for the original channel block port to wait before restoring the channel block and sending a R-APS Okay (node recovered remove block) message on receiving the R-APS recovery messages to allow the network to become stable. Range is 1-12 minutes. Default is 5 minutes.
West Logical Port	Port name. ***Note:*** Before configuring the west and east logical ports for a logical ring, disable RSTP and MSTP for the ports by using these commands: rstp disable port <PortName> and mstp disable port <PortName>

Copyright© 2022 Ciena® Corporation

Table 18-1
Attributes of logical rings

Attribute	Description
East Logical Port	Port name. ***Note:*** Before configuring the west and east logical ports for a logical ring, disable RSTP and MSTP for the ports by using these commands: rstp disable port <PortName> and mstp disable port <PortName>
West Logical Port Hold off Time	Time to wait before initiating ERP switching to allow for a link failure to recover, if possible. Range is 1-10 seconds. Default is 0.
East Logical Port Hold off Time	Time to wait before initiating ERP switching to allow for a link failure to recover, if possible. Range is 1-10 seconds. Default is 0.
West Logical Port CFM Service Name	Name of the CFM service associated with the MEP for the west logical port. Default is blank.
East Logical Port CFM Service Name	Name of the CFM service associated with the MEP for the east logical port. Default is blank.

Note 1: CFM must be used for rings which span transit (ERP switching unaware) switches or routers to account for all ports involved in the virtual ring. The CFM service can be configured for the data VLAN associated with the virtual ring. For details about configuring CFM services, see *39XX/51XX Service Delivery, Aggregation and Virtualization Advanced OAM Configuration*.

Note 2: Once a CFM service is associated with a port on a logical ring, a failure event from CFM against this port affects all virtual rings using this port as a ring port, even virtual rings that are children of logical rings that do not have this CFM service explicitly provisioned against them.

Note 3: For interoperability with other switches, such as the 5410, a separate VLAN is recommended for use with the CFM service monitoring the physical link.

CAUTION
CFM services msut not be associated with any R-APS channel VIDs.

Copyright© 2022 Ciena® Corporation

CAUTION
For CFM services used to monitor ring links, the only valid alarm-priority is 3 (default). Otherwise, after a failure, the ring never converges back to normal.

CAUTION
Risk of traffic flow blockage
When choosing the terminating port for a sub-ring as an east-port-termination or west-port-termination sub-ring port, make sure that the other corresponding logical-ring port is part of some other virtual ring (a major ring or another sub-ring). Otherwise, traffic may get blocked on the logical-ring port which is not included in any virtual ring.

The maximum number of logical rings supported for a device is equal to the number of its physical ports minus 1.

After creating a logical ring, you can modify the attributes that were specified at creation. If the logical ports are configured as MEPs, you can add their associated CFM service to trigger a switchover from faults detected by CFM.

You can:

- Create a logical ring
- Delete a logical ring. Before you can delete a logical ring, delete any virtual rings and any CFM services associated with it.
- Set logical ring attributes
- Restore logical ring attributes to default
- Force a port block, which forces a ring switch-over to block a non-RPL port and allow traffic to flow over the RPL port.
- Clear a port block, which clears a block from a port that has been forced into the block state and from ports blocked by ring protection.

Copyright© 2022 Ciena® Corporation

Step	Action

To create a logical ring

1 Create a logical ring:

```
ring-protection logical-ring create logical-ring-name
<logical-ring-name> {ring-id <NUMBER: 1-255>} {west-port
<PortNameList>} {east-port <PortNameList>} [guard-time
<MILLISECONDS>] [WTR <MINUTES>]
```

where

logical-ring-name <logical-ring-name>	is a unique alpha-numeric string to identify the logical ring.
ring-id <NUMBER: 1-255>	is a numeric identifier for the logical ring. SAOS supports "stacked" G.8032 logical rings with a common logical ring ID. This means: • SAOS allows the user to create more than one logical ring with the same logical ring ID. All virtual rings supported by a given node (no matter what ring ports are being used) are uniquely identified by the <logicalRingID, raps-VID> pair. • All virtual rings sharing a common logical ring port can have the same logical ring identifier (that is, <ringID>), and thus are uniquely identified by the <ringID, raps-VID> pair. For example, if the <ringID> is the same (for example, has a value of 1) across all logical rings, then the <raps-VID> needs to be unique.
west-port <Port list>	is the port identifier of the west port. • when a single link is used to transport traffic to the other switch, specify the physical port • when a link aggregation is used to transport traffic to the other switch, specify the logical port
east-port <Port list>	is the identifier of the east port. • when a single link is used to transport traffic to the other switch, specify the physical port. • when a link aggregation is used to transport traffic to the other switch, specify the logical port.
guard-time <MILLISECONDS>	is the guard time, which is used to prevent devices from acting on outdated R-APS messages and prevents the possibility of forming a closed loop. The default value is 500ms.
WTR <MINUTES>	is the wait to restore timer, which is used when the network is recovering from a signal failure to prevent multiple protection switching operations. The default value is 5 minutes.

Copyright© 2022 Ciena® Corporation

To delete a logical ring

2 Delete a logical ring:

```
ring-protection logical-ring delete logical-ring-name
<logical-ring-name>
```

where

logical-ring-name <logical-ring-name>	is the logical ring to delete.

To modify a logical ring

3 Modify logical ring attributes:

```
ring-protection logical-ring set ring <ring> [guard-time
<MILLISECONDS: 10-2000>] [wtr <MINUTES>] [west-port-
holdoff-time <MILLISECONDS: 0-10000>] [east-port-
holdoff-time <MILLISECONDS: 0-10000>] [west-port-cfm-
service <CFM Service Name>] [east-port-cfm-service <CFM
Service Name>]
```

where

ring <ring>	is the logical ring.
guard-time <MILLISECONDS: 10-2000>	is the guard time, which is used to prevent devices from acting on outdated R-APS messages and prevents the possibility of forming a closed loop. The default value is 500ms.
wtr <MINUTES>	is the wait to restore timer, which is used when the network is recovering from a signal failure to prevent multiple protection switching operations. The default value is 5 minutes.
west-port-holdoff-time <MILLISECONDS: 0-10000>	is the holdoff timer for the west port. The holdoff timer is used to coordinate timing of protection switches at multiple layers.
east-port-holdoff-time <MILLISECONDS: 0-10000>	is the holdoff timer for the east port. The holdoff timer is used to coordinate timing of protection switches at multiple layers.
west-port-cfm-service <CFM Service Name>	is the name of the CFM service configured for the west port logical ring links.
east-port-cfm-service <CFM Service Name>	is the name of the CFM service configured for the east port logical ring links.

Copyright© 2022 Ciena® Corporation

To restore logical ring attributes to default

4 Restore logical ring attributes to default:

```
ring-protection logical-ring unset ring <ring> {[guard-
time], [WTR], [west-port-holdoff-time] [east-port-
holdoff-time] [west-port-cfm-service-name] [east-port-
cfm-service-name]}
```

To force a port block

5 Force a port block:

```
ring-protection logical-ring force ring <ring> [west-
port] [east-port]
```

where

ring <ring>	is the logical ring.
[west-port]	is the west port.
[east-port]	is the east port.

To clear a port block

6 Clear a port block:

```
ring-protection logical-ring clear ring <ring> [west-
port] [east-port]
```

where

ring <ring>	is the logical ring.
[west-port]	is the west port.
[east-port]	is the east port.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example creates a logical ring.

```
ring-protection logical-ring create logical-ring-name lr1 ring-id 1 west-port
1 east-port 3
```

This example deletes the logical ring with the ring name lr1.

```
ring-protection logical-ring delete logical-ring-name lr1
```

This example sets logical ring attributes.

```
ring-protection logical-ring set ring lr1 guard-time 1000 east-port-cfm-
service-name VLAN#1000
```

This example restores logical ring attributes to default:

```
ring-protection logical-ring unset ring lr1 guard-time east-port-cfm-service-
name
```

This example forces a port block.

```
ring-protection logical-ring force ring lr1 west-port
```

This example clears the west-port port block on the logical ring lr1.

```
ring-protection logical-ring clear ring lr1 west-port
```


Copyright© 2022 Ciena® Corporation

Procedure 18-4
Configuring virtual rings

Create virtual rings to optimize usage of all links in a logical ring. You can create multiple virtual rings on the same logical ring by using different RPL owners.

For this reason, it is recommended that the RPL owner not be assigned to an interconnection node.

A sub-ring is a partial virtual ring, that is, a sub-ring does not form a closed loop. A sub-ring connects to two nodes of the major ring. This provides a dual-homed interconnection to the major ring in case of a failure at one of the interconnection nodes.

CAUTION
Risk of traffic flow blockage
When choosing the terminating port for a sub-ring as an east-port-termination or west-port-termination sub-ring port, make sure that the other corresponding logical-ring port is part of some other virtual ring (a major ring or another sub-ring). Otherwise, traffic may get blocked on the logical-ring port which is not included in any virtual ring.

After creating a virtual ring, you can modify the attributes that were specified when the virtual ring was created. The virtual ring inherits the attributes of the logical ring when it is created.

Additionally, set one of the logical ports to be RPL for the ring.

CAUTION
Risk of broadcast storms
Any VLANs that have not been added to a virtual ring are unprotected, meaning there is no mechanism to prevent broadcast storms.

CAUTION
Risk of traffic flow blockage
In sub-ring implementations, do not assign the RPL owner to an interconnection node. If the RPL owner is assigned to an interconnection node, faults on the sub-ring can prevent traffic flow to a node between the sub-ring endpoints.

You can create up to 4 rings per logical ring with a device maximum of twice the number of logical rings.

Copyright© 2022 Ciena® Corporation

Adding a VLAN to a virtual ring

You can associate one or more VLANs with a virtual ring. This VLAN can identify an enterprise VLAN defined by creating a VLAN and adding ports to it or it can be the provider VLAN (SVID) associated with the virtual circuit of a virtual switch for Q-in-Q encapsulation. Before adding a VLAN, make sure that the east and west logical ports are members of the VLAN.

Note 1: If more than one Virtual Ring is configured in a network, group VLANs together and maintain them in separate virtual rings throughout the entire network.

Note 2: The same VLAN cannot exist in both the G.8032 and MSTP realm.

Do not load-share VLANs across multiple virtual rings where there is a shared/interconnection node that also connects to other virtual ring(s) and the same VLANs are assigned to a single virtual ring.

The reason for this is that any VLAN or set of VLANs can belong only to a single "channel block group" even if a VLAN is assigned to several virtual rings with different channel block assignments at the interconnection node.

Modifying virtual ring attributes

After creating a virtual ring, you can modify the attributes that were specified at creation, and set one of the logical ports to be the RPL for the ring. Also, you can easily reset the attributes to their default settings.

You can:

- create a virtual ring
- delete a virtual ring. Before you can delete a virtual ring, remove any VLANs associated with it.
- add a VLAN to a virtual ring
- remove a VLAN from a virtual ring
- set virtual ring attributes
- restore virtual ring attributes to default
- move a VLAN from one virtual ring to another virtual ring

Copyright© 2022 Ciena® Corporation

Step	Action

To create a virtual ring

1 Create a virtual ring:

```
ring-protection virtual-ring create virtual-ring-name
<virtual-ring-name> {logical-ring <Logical ring name>}
{raps-vid <NUMBER>} [mel <NUMBER: 0-7>] [sub-ring <no-
port-termination|west-port-termination|east-port-
termination>]
```

where

virtual-ring-name <virtual-ring-name>	is the unique alpha-numeric string to identify the virtual ring.
logical-ring <Logical ring name>	is the associated logical ring name.
raps-vid <NUMBER>	is the VLAN-ID assigned to the R-APS channel. This VID number cannot overlap with any data VLANs configured on the device. On creation of the virtual ring, the R-APS VID is entered in the VLAN table. It can be used for multiple virtual rings that are associated with different logical rings.
mel <NUMBER: 0-7>	is the maintenance entity group level. The default value is 0. In a CFM configuration, this attribute is called the MD level.
sub-ring <no-port-termination| west-port-termination| east-port-termination>	is the sub-ring. In sub-ring configurations, identifies whether the sub-ring terminates on the east ring port, west ring port, or does not terminate on this note. This setting is required for all nodes in the sub-ring.

To delete a virtual ring

2 Delete a virtual ring:

```
ring-protection virtual-ring delete ring <ring>
```

where

ring <ring>	is the name of the virtual ring that you want to delete.

Copyright© 2022 Ciena® Corporation

To add a VLAN to a virtual ring

3 Add a VLAN to a virtual ring:

```
ring-protection virtual-ring add ring <ring> vid <VLAN
list>
```

where

ring <ring>	is the virtual ring name.
vid <VLAN list>	is the VLAN ID of the VLAN that you want to add to the virtual ring.

To remove a VLAN from a virtual ring

4 Remove a VLAN from a virtual ring:

```
ring-protection virtual-ring remove ring <String> vid
<VlanList>
```

where

ring <ring>	is the name of the virtual ring that you want to remove the VLAN from.
vid <VLAN list>	is the VLAN ID of the VLAN that you want to remove.

Copyright© 2022 Ciena® Corporation

To modify a virtual ring

5 Set virtual ring attributes:

```
ring-protection virtual-ring set ring <ring> {[mel
<NUMBER: 0-7>], [revertive <off|on>], [west-port-rpl
<owner>] [east-port-rpl <owner>]}
```

where

ring <ring>	is a unique alpha-numeric string to identify the virtual ring.
mel <NUMBER: 0-7>	is the maintenance entity group level. The default value is 0. In a CFM configuration, this attribute is called the MD level.
revertive <off|on>	determines whether the revertive attribute is on or off. • on: after a failure condition is cleared, the traffic channel unblocks the ports for the failed link and reverts to the configured RPL • off: after a failure condition is cleared, the traffic channel continues to use the previously-failed node as the channel block, and traffic continues over the configured RPL. If a second failure occurs, the recovered link unblocks its ports When revertive is set to “off” recovering from a failure, the virtual ring does not migrate back to the “OK” state. It stays in the “Recovering” state until the ring is manually reverted. Manual reversion is done by using the logical ring “clear” command on the RPL owned port and migrates the block back to the RPL owned port, thereby changing the virtual ring to the “OK” state from the “Recovering” state.
west-port-rpl <owner>	is the west RPL port. Only one port of all devices in the virtual ring can be the RPL. By default, neither port is set to be the RPL.
east-port-rpl the <owner>	is the east RPL port. Only one port of all devices in the virtual ring can be the RPL. By default, neither port is set to be the RPL.

Copyright© 2022 Ciena® Corporation

To restore virtual ring attributes to default

6 Restore virtual ring attributes to default:

```
ring-protection virtual-ring unset ring <ring> {[mel],
[revertive], [west-port-rpl], [east-port-rpl]}
```

where

ring <ring>	is the virtual ring.
[mel]	is the maintenance entity group level.
revertive	indicates revertive.
[west-port-rpl]	is the RPL.
[east-port-rpl]	is the RPL.

To move a VLAN from one virtual ring to another virtual ring

7 Break the ring or rings to prevent a loop by disabling the port that the ring or rings terminate on:

```
port disable port <port>
```

where

port <port>	is the port that the ring or rings terminate on.

8 Remove the VLAN from all current major virtual rings and virtual sub-rings:

```
ring-protection virtual-ring remove ring <String> vid
<VlanList>
```

where

ring <ring>	is the name of the virtual ring that you want to remove the VLAN from.
vid <VLAN list>	is the VLAN ID of the VLAN that you want to remove.

9 Add the VLAN to the other selected virtual ring or rings:

```
ring-protection virtual-ring add ring <ring> vid <VLAN
list>
```

where

ring <ring>	is the name of the virtual ring that you want to delete.
vid <VLAN list>	is the VLAN ID of the VLAN that you want to add to the virtual ring.

10 Re-enable the downed port:

```
port enable port <port>
```

where

port <port>	is the port that the ring or rings terminate on.

Copyright© 2022 Ciena® Corporation

11 Verify that the RPL owner is not learning any MACs on the VLANs that must be blocked:

```
flow mac-addr show vlan <vlan> [mac <MAC address:
XX:XX:XX:XX:XX:XX>] [table-count]
```

where

vlan <vlan>	is the VLAN list.
mac <MAC address: XX:XX:XX:XX:XX:XX>	is the MAC address.
table-count	displays the table count only.

```
flow mac-addr show port <port>
```

where

port <port>	is the port number.

—end—

Example

This example creates a virtual ring with the name vr1. The associated logical ring is lr1. The dedicated R-APS VID is 100.

```
ring-protection virtual-ring create virtual-ring-name vr1 logical-ring lr1
raps-vid 100
```

This example deletes the virtual ring named vr1.

```
ring-protection virtual-ring delete ring vr1
```

This example adds an enterprise VLAN to a virtual ring.

```
vlan create vlan 1000
ring-protection virtual-ring add ring vr1 vid 1000
```

This example adds a provider VLAN (SVID) to a virtual ring.

```
vlan create vlan 1000
virtual-circuit ethernet create vc vc1 vlan 1000
virtual-switch add reserved-vlan 4000
virtual-switch ethernet create vs vs1 vc vc1
ring-protection virtual-ring add ring vr1 vid 1000
```

This example removes the VLAN with the vid 1000 from the virtual ring named vr1.

```
ring-protection virtual-ring remove ring vr1 vid 1000
```

This example sets virtual ring attributes.

```
ring-protection virtual-ring set ring vr1 west-port-rpl owner
```


Copyright© 2022 Ciena® Corporation

This example restores virtual ring attributes to their default values.

```
ring-protection virtual-ring unset ring vr1 west-port-rpl
```


Copyright© 2022 Ciena® Corporation

Procedure 18-5
Verifying ERP switching operation

Verify ERP switching operation by forcing a ring switchover to block a non-RPL port, which allows traffic to flow over the RPL port.

After you have verified that the traffic flows correctly, clear the block.

Step	Action

To force a port block

1 Force a port block:

```
ring-protection logical-ring force ring <LogicalRingName>
[west-port|east-port]
```

where

force ring <LogicalRingName>	is the name of the logical ring.
[west-port|east-port]	selects the west or east port.

2 Verify traffic flow:

```
ring-protection show
```

To clear a port block

3 Clear a port block:

```
ring-protection logical-ring clear <LogicalRingName>
[west-port|east-port]
```

where

force ring <LogicalRingName>	is the name of the logical ring.
[west-port|east-port]	selects the west or east port.

—end—

Example

This example forces a port block.

```
ring-protection logical-ring force ring 1-ring west-port
```

This example clears a port block.

```
ring-protection logical-ring clear ring 1-ring west-port
```


Copyright© 2022 Ciena® Corporation

Procedure 18-6
Displaying information for logical rings

Display information for logical rings to verify configuration. You can display:

- a list of logical rings with the number of virtual rings for each logical ring
- details for a specific logical ring

Step	Action

To display logical rings

1 Display logical rings:

```
ring-protection logical-ring show
```

To display details for a specific logical ring

2 Display details about a specific logical ring:

```
ring-protection logical-ring show [ring <RingName>]
```

where

show [ring <RingName>]	is the name of the logical ring.

—end—

Example

This example shows a list of logical rings.

```
ring-protection logical-ring show

+----------------- RING-PROTECTION LOGICAL-RINGS -----------------+
| Name                              | Number of Virtual Rings       |
+-----------------------------------+-------------------------------+
| lr1                               | 1                             |
| lr2                               | 1                             |
+-----------------------------------+-------------------------------+
```


Copyright© 2022 Ciena® Corporation

This example shows details for a logical ring with a ring name of lr1.

```
ring-protection logical-ring show ring lr1

+----------------------- RING-PROTECTION LOGICAL-RING lr1 -----------------------+
| Parameter            | Value                                                   |
+----------------------+---------------------------------------------------------+
| Ring ID              | 1                                                       |
| Guard Time (msec)    | 1000                                                    |
| WTR (min)            | 5                                                       |
| No of Virtual Rings  | 1                                                       |
+------------------------- LOGICAL-RING RING-PORT INFO --------------------------+
| Parameter            | West Port                  | East Port                  |
+----------------------+----------------------------+----------------------------+
| Port                 | 1                          | 3                          |
| Holdoff Time (msec)  | 0                          | 0                          |
| Force Switch         | Off                        | Off                        |
| CFM Service          |                            | VLAN#1000                  |
+----------------------+----------------------------+----------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 18-7
Displaying virtual ring information

Display virtual ring information to verify configuration.

You can display a summary of configured virtual rings or details about a specific virtual ring. The summary includes a list of virtual rings along with their state. This table lists virtual ring states.

Table 18-2
Virtual ring states

State	Description
Ok	Indicates protection switching is not active and the block is on the RPL.
Protecting	Indicates there is an active fault on the ring.
Recovering	Indicates that there is no active fault on the ring and the block has not yet moved to the RPL port. In revertive mode, the “recovering” state lasts for the duration of the WTR, and then the block is moved back to the RPL. In non-revertive mode, the recovering state lasts until the block is cleared.

Details about a specific virtual ring include:

- configuration
- number of switch-overs
- up-time
- down-time
- ring port configuration
- status
- statistics

The Wait To Block (WTB) time shown as part of the details for a specific ring is not configurable. It is always five seconds greater than the configurable Guard Time. It is used exclusively for recovery from a ring-protection logical-ring clear command. When ring-protection logical-ring clear command is issued to recover the ring to its normal state, the WTB time is the waiting interval after which recovery is actually initiated.

Copyright© 2022 Ciena® Corporation

Step	Action

To display a summary of configured virtual rings

1 Display a summary:

```
ring-protection virtual-ring show [ring
<VirtualRingName>]
```

where

show [ring <VirtualRingName>] is the name of the virtual ring.

To display details of a virtual ring

2 Display details about a virtual ring:

```
ring-protection virtual-ring show [ring
<VirtualRingName>]
```

where

show [ring <VirtualRingName>] is the name of the virtual ring.

—end—

Example

This example displays a summary of configured virtual rings.

```
> ring-protection virtual-ring show

+---------------- RING-PROTECTION VIRTUAL-RINGS -----------------+
| name                              | State                        |
+-----------------------------------+------------------------------+
| vr1                               | Ok                           |
| vr2                               | Ok                           |
+-----------------------------------+------------------------------+
```


Copyright© 2022 Ciena® Corporation

This example shows ring-protection virtual-ring information for a virtual ring with the name vr1.

```
> ring-protection virtual-ring show ring vr1

+-------------------- RING-PROTECTION VIRTUAL-RING vr1 ----------------------+
| Parameter                          | Value                                  |
+------------------------------------+----------------------------------------+
| Ring Type                          | Major-ring                             |
| RAPS VID                           | 100                                    |
| Logical Ring                       | lr1                                    |
| MEL                                | 0                                      |
| State                              | Ok                                     |
| Status                             | Ok                                     |
| Revertive                          | On                                     |
| Guard Time                         | 500                                    |
| WTR (Wait To Restore) Time (min)   | 5                                      |
| WTB (Wait To Block) Time (msec)    | 5500                                   |
| No of Switchovers                  | 0                                      |
| Uptime From Last Failure           | 22m 0s                                 |
| Total Down Time                    | 0s                                     |
+------------------------------------+----------------------------------------+
+----------------------- VIRTUAL-RING RING-PORT INFO ------------------------+
| Parameter                 | West Port                | East Port                |
+---------------------------+--------------------------+--------------------------+
| Port                      | 1                        | 3                        |
| Sub-ring Termination      | Not a Sub-ring Port      | Not a Sub-ring Port      |
| Holdoff Time (msec)       | 0                        | 0                        |
| Force Switch              | Off                      | Off                      |
| RPL                       | None                     | None                     |
| State                     | Forwarding               | Forwarding               |
| Status                    | Ok                       | Ok                       |
+---------------------------+--------------------------+--------------------------+
+------------------- VIRTUAL-RING RING-PORT STATISTICS ---------------------+
|                           |       West Port        |       East Port        |
| R-APS Frame Type          |  Rx Pkts  |  Tx Pkts   |  Rx Pkts  |  Tx Pkts   |
+---------------------------+-----------+------------+-----------+------------+
| No Request (NR)           |        12 |         31 |        12 |         31 |
| Signal Fail (SF)          |         0 |         10 |         0 |         10 |
| Forced Switch (FS)        |         0 |          0 |         0 |          0 |
| NR, RPL Blocked (RB)      |         0 |        266 |         0 |        266 |
+---------------------------+-----------+------------+-----------+------------+
+----------------------- VIRTUAL-RING MEMBER VLANS ------------------------+
| 1000                                                                      |
+---------------------------------------------------------------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 18-8
Displaying ring protection

Display ring protection to verify the configuration. The Ring-Protection Info table displays:

- state of ring protection
- node ID
- EtherType
- number of configured logical rings for the device

Step	Action
1	Display ring protection: `ring-protection show`

—end—

Example

This example shows sample output for the ring-protection show command.

```
> ring-protection show

+--------------- RING-PROTECTION INFO ----------------+
| Parameter                    | Value                 |
+------------------------------+-----------------------+
| State                        | Enabled               |
| Node ID                      | 00:02:5A:01:B0:AA     |
| Ether Type                   | 0x8100                |
| Number Of Logical Rings      | 1                     |
+------------------------------+-----------------------+
```


Copyright© 2022 Ciena® Corporation

Procedure 18-9
Displaying ERP switching in the configuration file

Display ERP switching in the configuration file as part of troubleshooting the ERP configuration.

ERP switching configuration is saved to the configuration file in the RING PROTECTION CONFIG: and RING PROTECTION CFM SERVICE CONFIG: sections.

```
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
! RING PROTECTION CONFIG:
!
ring-protection logical-ring create logical-ring-name lr1 ring-id 1 west-port
1 east-port 3
ring-protection logical-ring create logical-ring-name lr2 ring-id 2 west-port
3 east-port 2
ring-protection virtual-ring create virtual-ring-name vr1 logical-ring lr1
raps-vid 100
ring-protection virtual-ring add ring vr1 vid 1000
ring-protection virtual-ring create virtual-ring-name vr2 logical-ring lr2
raps-vid 102 sub-ring east-port-termination
ring-protection virtual-ring add ring vr2 vid 1000
!...
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!

! RING PROTECTION CFM SERVICE CONFIG:

!

ring-protection logical-ring set ring lr1 east-port-cfm-service-name
VLAN#1000

!

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
```


Copyright© 2022 Ciena® Corporation

Procedure 18-10
Checking configuration for a ring not forwarding traffic

Check the configuration for a ring not forwarding traffic when troubleshooting.

After configuring the logical ring(s) and virtual ring(s), if traffic is not forwarding (as indicated by the state of the virtual ring ports) for all ports in the ring other than the RPL port, make sure to check that:

- Logical ring Ring Identifier (ring-id) matches for each node.
- Virtual ring R-APS VID (raps-vid) matches for each node.
- Virtual ring VLAN (vlan) matches for each node.
- One RPL owner is defined for the entire virtual ring (on one node only).
- Virtual ring is in the normal “Ok” state for each node.
- R-APS messages are transmitted and received for each node.

Copyright© 2022 Ciena® Corporation

Procedure 18-11
Configuration example

In this example, nodes A, B, C, and D are connected in a loop to create the major ring with R-APS VID 100. Nodes B, D, E and F are also members of the sub-ring with R-APS VID 102. It is assumed that the configuration is already in place, either for enterprise VLAN 1000 or for the provider VLAN 1000 associated with the virtual circuit of a virtual switch for Q-in-Q encapsulation.

Figure 18-2
Example with logical ring, major ring, and sub-ring

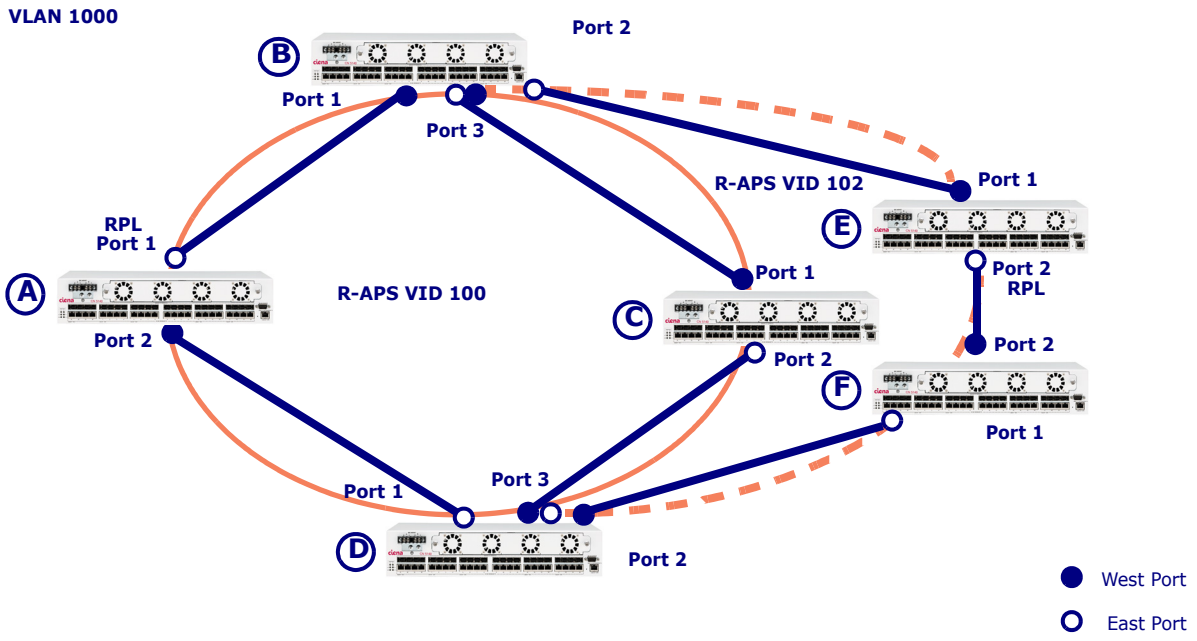

Node A

Disable the spanning tree protocol for the west and east logical ports.

```
rstp disable port 1,2
mstp disable port 1,2
```

Create the logical ring.

```
ring-protection logical-ring create logical-ring-name lr1 ring-
id 1 west-port 2 east-port 1
```

Create the virtual ring using a R-APS VID different from the enterprise or provider VLAN.

Copyright© 2022 Ciena® Corporation

```
ring-protection virtual-ring create virtual-ring-name vr1
logical-ring lr1 raps-vid 100
```

Add the enterprise or provider VLAN to the virtual ring.

```
ring-protection virtual-ring add ring vr1 vid 1000
```

Set the RPL port for the virtual ring.

```
ring-protection virtual-ring set ring vr1 east-port-rpl owner
revertive on
```

Validate the configuration and state (optional).

```
ring-protection logical-ring show ring lr1
ring-protection virtual-ring show ring vr1
```

Node B

Disable the spanning tree protocol for the west and east logical ports.

```
rstp disable port 1,2,3
mstp disable port 1,2,3
```

Create the logical ring for the major ring.

```
ring-protection logical-ring create logical-ring-name lr1 ring-
id 1 west-port 1 east-port 3
```

Create the logical ring for the sub-ring.

```
ring-protection logical-ring create logical-ring-name lr2 ring-
id 2 west-port 3 east-port 2
```

Create the major virtual ring using a R-APS VID different from the enterprise or provider VLAN.

```
ring-protection virtual-ring create virtual-ring-name vr1
logical-ring lr1 raps-vid 100
```

Add the enterprise or provider VLAN to the major virtual ring.

```
ring-protection virtual-ring add ring vr1 vid 1000
```

Create the virtual ring for the sub-ring using a R-APS VID different from the enterprise or provider VLAN and the major virtual ring.

```
ring-protection virtual-ring create virtual-ring-name vr2
logical-ring lr2 raps-vid 102 sub-ring east-port-termination
```


Copyright© 2022 Ciena® Corporation

Add the VLAN to the sub-ring.

```
ring-protection virtual-ring add ring vr2 vid 1000
```

Validate the configuration and state (optional).

```
ring-protection logical-ring show ring lr1
ring-protection logical-ring show ring lr2
ring-protection virtual-ring show ring vr1
ring-protection virtual-ring show ring vr2
```

Node C

Disable the spanning tree protocol for the west and east logical ports.

```
rstp disable port 1,2
mstp disable port 1,2
```

Create the logical ring.

```
ring-protection logical-ring create logical-ring-name lr1 ring-
id 1 west-port 1 east-port 2
```

Create the virtual ring using a R-APS VID different from the enterprise or provider VLAN.

```
ring-protection virtual-ring create virtual-ring-name vr1
logical-ring lr1 raps-vid 100
```

Add the VLAN to the virtual ring.

```
ring-protection virtual-ring add ring vr1 vid 1000
```

Validate the configuration and state (optional).

```
ring-protection logical-ring show ring lr1
ring-protection virtual-ring show ring vr1
```

Node D

Disable the spanning tree protocol for the west and east logical ports.

```
rstp disable port 1,2,3
mstp disable port 1,2,3
```

Create the logical ring for the major ring.

```
ring-protection logical-ring create logical-ring-name lr1 ring-
id 1 west-port 3 east-port 1
```


Copyright© 2022 Ciena® Corporation

Create the logical ring for the sub-ring.

```
ring-protection logical-ring create logical-ring-name lr2 ring-
id 2 west-port 2 east-port 3
```

Create the major virtual ring using a R-APS VID different from the enterprise or provider VLAN.

```
ring-protection virtual-ring create virtual-ring-name vr1
logical-ring lr1 raps-vid 100
```

Add the enterprise or provider VLAN to the major virtual ring.

```
ring-protection virtual-ring add ring vr1 vid 1000
```

Create the virtual ring for the sub-ring using a R-APS VID different from the enterprise or provider VLAN and the major virtual ring.

```
ring-protection virtual-ring create virtual-ring-name vr2
logical-ring lr2 raps-vid 102 sub-ring west-port-termination
```

Add the VLAN to the sub-ring.

```
ring-protection virtual-ring add ring vr2 vid 1000
```

Validate the configuration and state (optional).

```
ring-protection logical-ring show ring lr1
ring-protection logical-ring show ring lr2
ring-protection virtual-ring show ring vr1
ring-protection virtual-ring show ring vr2
```

Node E

Disable the spanning tree protocol for the west and east logical ports.

```
rstp disable port 1,2
rstp disable port 1,2
```

Create the logical ring.

```
ring-protection logical-ring create logical-ring-name lr2 ring-
id 2 west-port 1 east-port 2
```

Create the virtual ring for the sub-ring using a R-APS VID different from the enterprise or provider VLAN.

```
ring-protection virtual-ring create virtual-ring-name vr2
logical-ring lr2 raps-vid 102 sub-ring no-port-termination
```


Copyright© 2022 Ciena® Corporation

Add the VLAN to the virtual ring.

```
ring-protection virtual-ring add ring vr2 vid 1000
```

Set the RPL port for the virtual ring.

```
ring-protection virtual-ring set ring vr2 east-port-rpl owner
revertive on
```

Validate the configuration and state (optional).

```
ring-protection logical-ring show ring lr2
ring-protection virtual-ring show ring vr2
```

Node F

Disable the spanning tree protocol for the west and east logical ports.

```
rstp disable port 1,2
mstp disable port 1,2
```

Create the logical ring.

```
ring-protection logical-ring create logical-ring-name lr2 ring-
id 2 west-port 2 east-port 1
```

Create the virtual ring using a R-APS VID different from the enterprise or provider VLAN.

```
ring-protection virtual-ring create virtual-ring-name vr2
logical-ring lr2 raps-vid 102 sub-ring no-port-termination
```

Add the VLAN to the virtual ring.

```
ring-protection virtual-ring add ring vr2 vid 1000
```

Validate the configuration and state (optional).

```
ring-protection logical-ring show ring lr2
ring-protection virtual-ring show ring vr2
```


Copyright© 2022 Ciena® Corporation

Virtual Link Loss Indication configuration

Virtual Link Loss Indication (VLLI) enables a 39XX/51XX switch at the user/network interface (UNI) port of an EVC to signal to the connected UNI of the customer premise equipment (CPE) if one of these conditions occurs:

- the remote UNI attachment interface has failed or recovered
- the EVC connection has failed or recovered

In case of failure, the affected UNI port can be disabled and traffic rerouted until the problem is resolved.

When provisioning VLLI for a service, you can specify the triggering condition and the response in case of service failure and subsequent service restoration.

This section discusses these topics:

- "VLLI signaling" on page 19-1
- "VLLI architecture" on page 19-3
- "Y-cable protection use case" on page 19-5
- "Procedures" on page 19-7

VLLI signaling

VLLI monitors the continuity of the EVC by means of

- CCMs
- physical port status at local endpoints

CCMs contain an interface status TLV that conveys the status of the interface, that is, up or down, between the local and remote interfaces. CCM interface status and CCM timeout messages are used to convey remote interface status information.

This figure shows an example EVC between a Service Delivery Switch (SDS) and a 39XX/51XX switch.

Copyright© 2022 Ciena® Corporation

Figure 19-1
EVC between SDS and 39XX/51XX switch

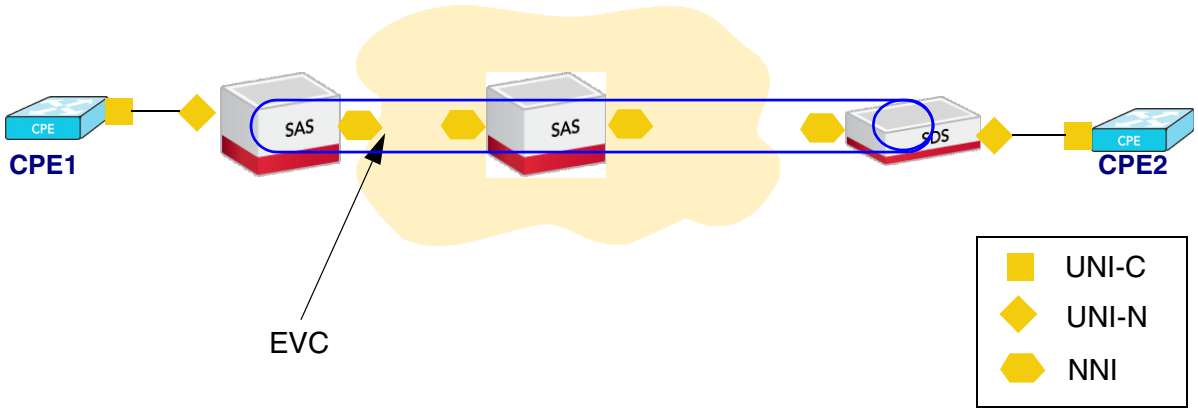

This figure shows a fiber cut or equipment failure between CPE2 and the SDS. VLLI signals CPE1.

Figure 19-2
Fiber cut or equipment failure between CPE2 and SDS

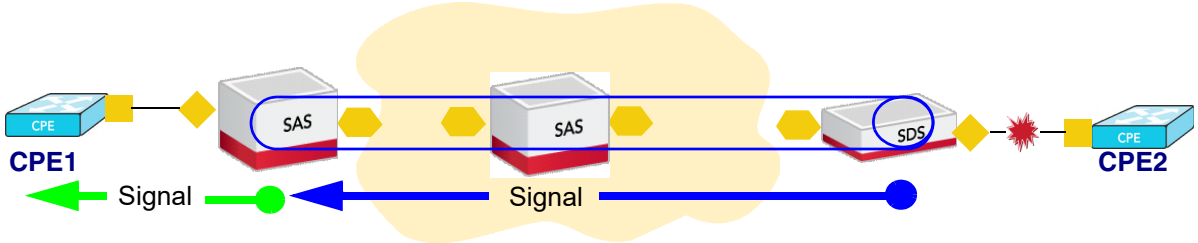

This figure shows a fiber cut or equipment failure between two 39XX/51XX switches. VLLI signals CPE1 and CPE2.

Figure 19-3
Fiber cut or equipment failure between two 39XX/51XX switches

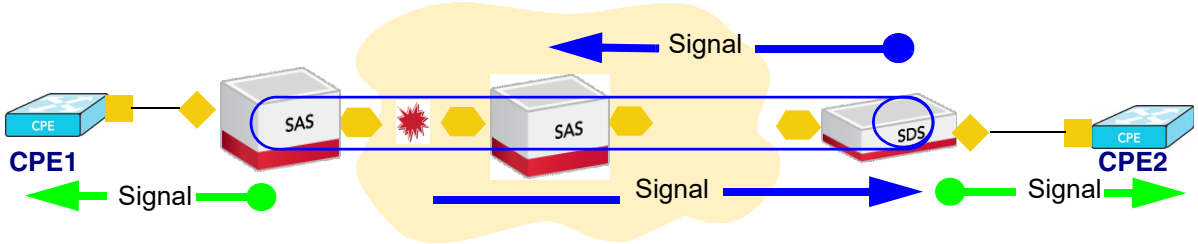

Copyright© 2022 Ciena® Corporation

VLLI architecture

An instance group includes multiple instances, to a maximum of 256. An instance represents an application service, for example, a CFM service. An instance corresponds to source or destination. One instance triggers the associated action to be performed on the other instances in the same instance group. Actions are

- ccm-stop, which stops the transmission of CCMs
- ccm-resume, which restarts the transmission of CCMs
- port-shut, which disables the port
- port-unshut, which restores the port
- mac-flush, which flushes all the MAC addresses learned on the port
- vc-mac-withdraw, which sends a mac-withdraw on the PW
- pw-status-fault, which sends an AC:Tx:Rx fault ot the mpls l2VpnMgr
- pw-status-clear, which sends an ACT:Tx:Rx clear to the mpls L2VpnMgr

Each instance links to an action group. An action group is a set of actions.

This figure shows a sample instance group and action group configuration.

Figure 19-4
Sample instance group and action group configuration

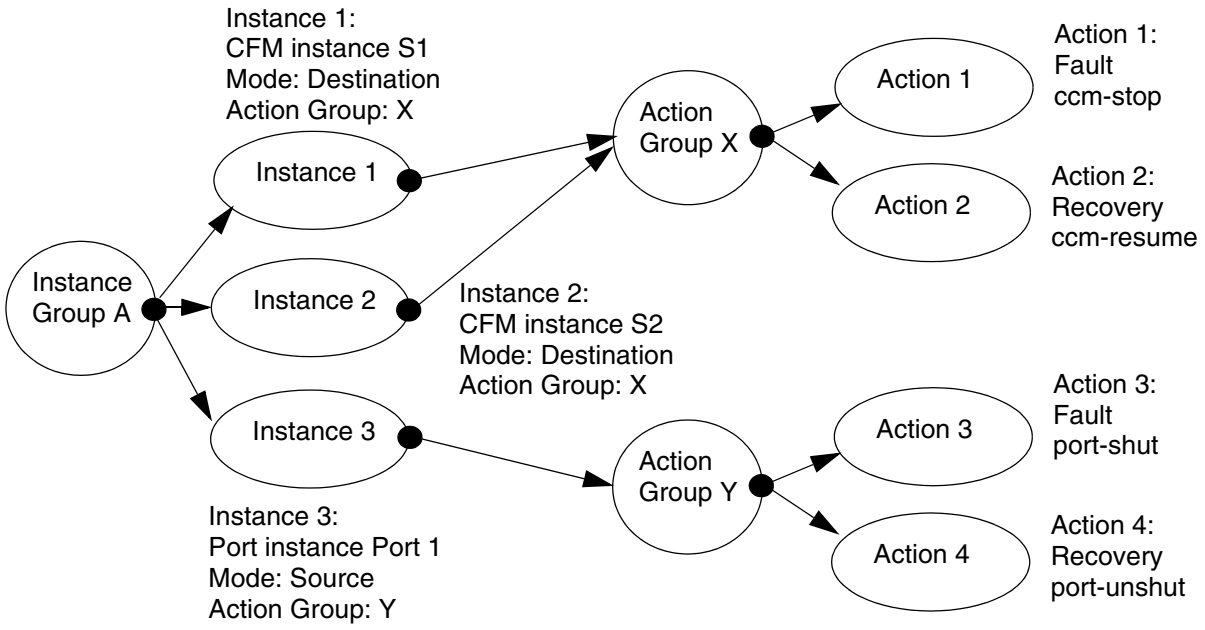

Copyright© 2022 Ciena® Corporation

Assume that instance 1 and instance 2 detect CCM loss. VLLI checks that all source destinations have failed before triggering the action on instance 3, which is to shut the port.

> ***Note:*** The VLLI instance-group delete command does not return an error if vc-instances are added to instance-group. Therefore, when the configuration has multiple vc-instances in an instance group, the user is expected to remove vc-instances from the instance-group, rather than deleting the instance-group altogether.

This table lists attributes for action groups.

Table 19-1
Attributes for action groups

Attribute	Description
action-group <action-group>	is the virtual-link-loss-indication action group.
type <fault | recovery>	is the type.
action <port-shut | port-unshut | ccm-stop | ccm-resume | mac-flush | vc-mac-withdraw | pw-status-fault | pw-status-clear>	Is the application action name.
precedence <NUMBER: 1-8>	is the precedence, which defines the order of actions in an action group. Actions are executed based on this order.

This table lists attributes for instance groups.

Table 19-2
Attributes for instance groups

Attribute	Description
instance-group <instance-group>	is the virtual-link-loss-indication instance group name.
mode <source| destination>	is the mode.
cfm-instance <CFM Service Name>	is the CFM service.
port-instance <Port list>	is the source port list.

Copyright© 2022 Ciena® Corporation

Table 19-2
Attributes for instance groups

Attribute	Description
vc-instance	is the PW instance.
action-group <VLLI action>	is the action group name.
direction <unidirectional | bidirectional>]	is the direction. Valid values are: • unidirectional: event on source instance triggers a VLLI action on the designated destination instance • bidirectional: event on source instance can trigger a VLLI action on the destination instance AND event on destination instance can trigger a VLLI action on the source instance VLLI supports unidirectional and bidirectional operation on the same instance group.
trigger <all | any>	is the trigger condition for an instance group. Valid values are: • all: the switch detects all the same events from the source instance before triggering the target's action • any: the target's action is triggered if any source's event is detected

Y-cable protection use case

A Y-cable protection scheme employs two Y-cables, which are hardware combiner/splitters.

A signal injected into the stem of the Y is duplicated into both arms of the Y with 50 percent attenuation in each arm. Signals injected into the arms of the Y are summed into the stem of the Y.

On the operational side, if both Tx lasers at the split ends are turned on, and if the common end starts receiving signals from both of the split-end ports, the Y-cable topology may behave undeterministically and the topology may go down. Therefore, the Tx laser at one split end must be on, and the Tx laser at the other split end must be off. To achieve this, both split ends are controlled by VLLI.

Since a Y-cable protection scheme provides redundancy and not load balancing, only one of the arms should be injecting traffic towards the stem. It is assumed that Y-cable is deployed along with the standard compliant pluggable.

The following figure depicts a typical customer use case where one port of CPE-1 is connected with two Ciena nodes A and B through a Y-cable:

Copyright© 2022 Ciena® Corporation

Figure 19-5
Y-cable protection use case scenario

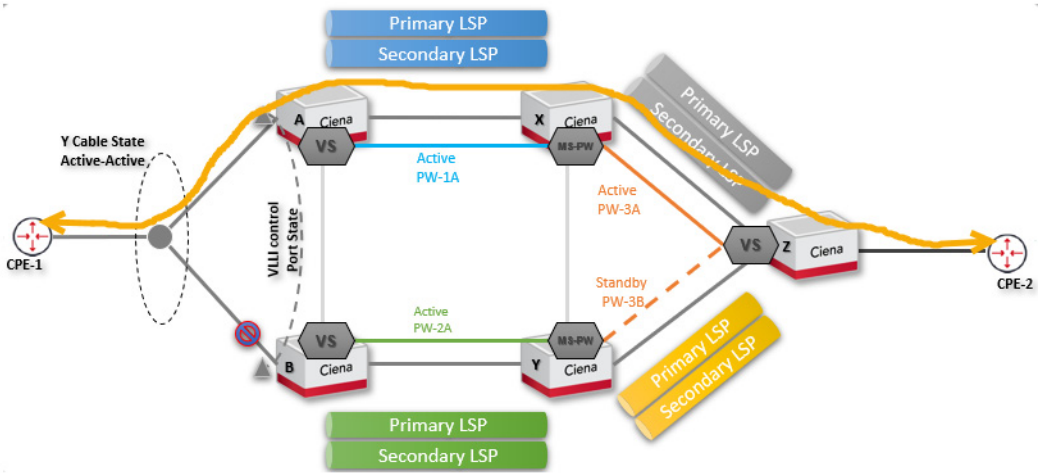

CPE-2 is connected to A and B through Ciena nodes X and Y through active/active pseudowire configuration.

For a description of the Y-cable feature, see *39XX/51XX Switches and Platforms Base Configuration*.

For procedures related to configuring the ports for Y-cable support, see the *Managing ports* chapter in *39XX/51XX Switches and Platforms Base Configuration*.

For descriptions of Y-cable failure and reversion scenarios, see "Y-cable protection failure and reversion scenarios".

Y-cable protection failure and reversion scenarios

Failure scenarios. In case of a fiber cut OR a port failure on the active node (which is carrying the traffic), the corresponding port on the standby node will come up and the node will assume the active role.

Dual node interconnect link (ICL) down. This would lead to a split brain condition and should never occur in the field. It is recommended that ICL CFM service should always be protected.

Automatic reversion. Automatic reversion is not supported.

Manual reversion. Assuming that the secondary node has become active and the primary node has taken the backup role, disable the y-protection port on the secondary node. That will cause the y-protection port on the primary node to become active. When the primary-node link comes up,

Copyright© 2022 Ciena® Corporation

administratively enable the port on the secondary again. That ensures that if the link through primary node goes down, the secondary node will be able to bring up the link.

Procedures

This section provides these procedures:

- “Configuring VLLI to shut down a port when CCM loss is detected - virtual switch” on page 19-8
- “Configuring VLLI to shut down a port when CCM loss is detected - VLAN” on page 19-13
- “Configuring VLLI between a UNI and an ENNI” on page 19-16
- “Configuring inverse VLLI with one 39XX/51XX switch and two PEs” on page 19-20
- “Configuring inverse VLLI with two 39XX/51XX switches and two PEs” on page 19-23
- “Configuring VLLI for MAC flush on an intersecting node” on page 19-26
- “Configuring MAC flush from the MPLS network ingress node to all other nodes” on page 19-28
- “Y-cable client protection using inverse VLLI between nodes B and A” on page 19-30
- “MPLS active/active protection state management using VLLI in a Y-cable topology” on page 19-33
- “MPLS active/standby protection state management using VLLI in a Y-cable topology” on page 19-42
- “Configuring VLLI to block a pseudowire” on page 19-49

Copyright© 2022 Ciena® Corporation

Procedure 19-1
Configuring VLLI to shut down a port when CCM loss is detected - virtual switch

Configure VLLI to automatically shut down a port when CCM loss is detected. The figures provided in this procedure show sample scenarios for shutting down a port when VLLI detects CCM loss. These scenarios are based on a CFM service running on a virtual switch.

This figure shows a fiber cut or equipment failure that impacts the connection between two 39XX/51XX switches.

Figure 19-6
Connection between 39XX/51XX switches lost due to fiber cut or equipment failure

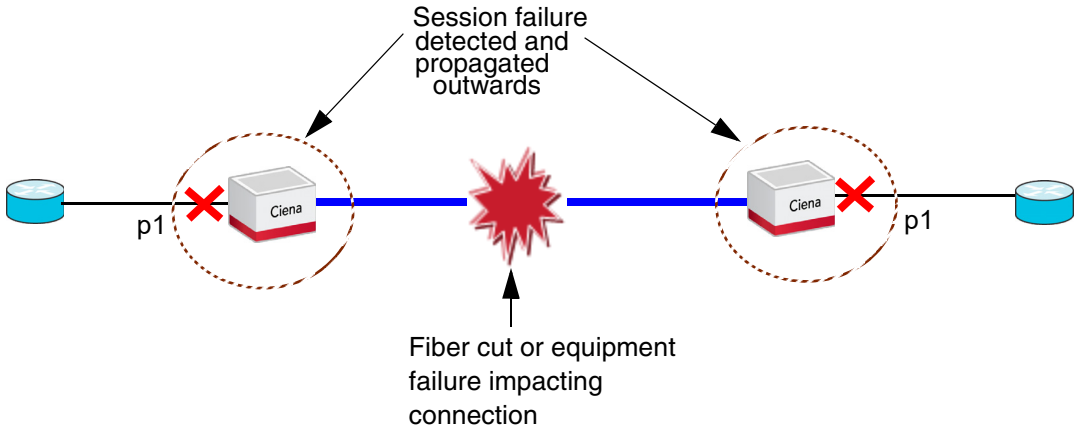

This figure shows a fiber cut or equipment failure that impacts the connection between a 39XX/51XX switch and a CPE.

Copyright© 2022 Ciena® Corporation

Figure 19-7
Connection between 39XX/51XX switch and CPE lost due to fiber cut or equipment failure

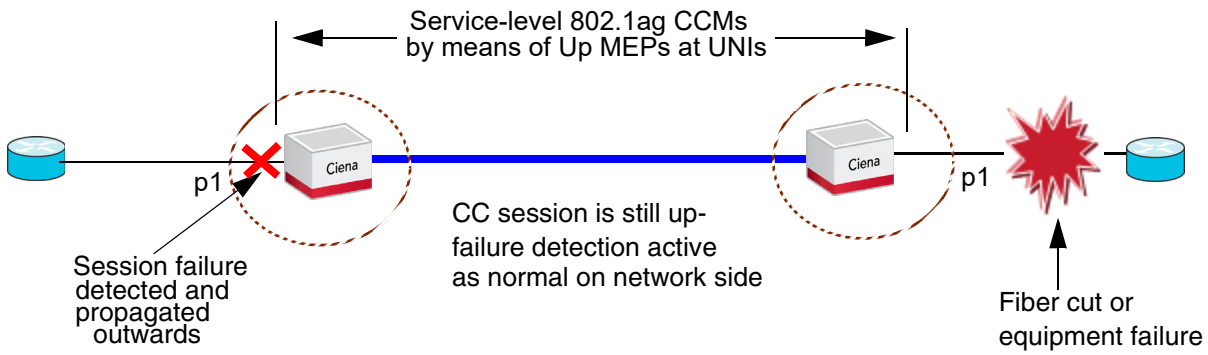

This figure shows a fiber cut or equipment failure that impacts the connection between two 39XX/51XX switches through a G.8032 ring.

Figure 19-8
Connection between 39XX/51XX switches through a G.8032 ring lost due to fiber cut or equipment failure

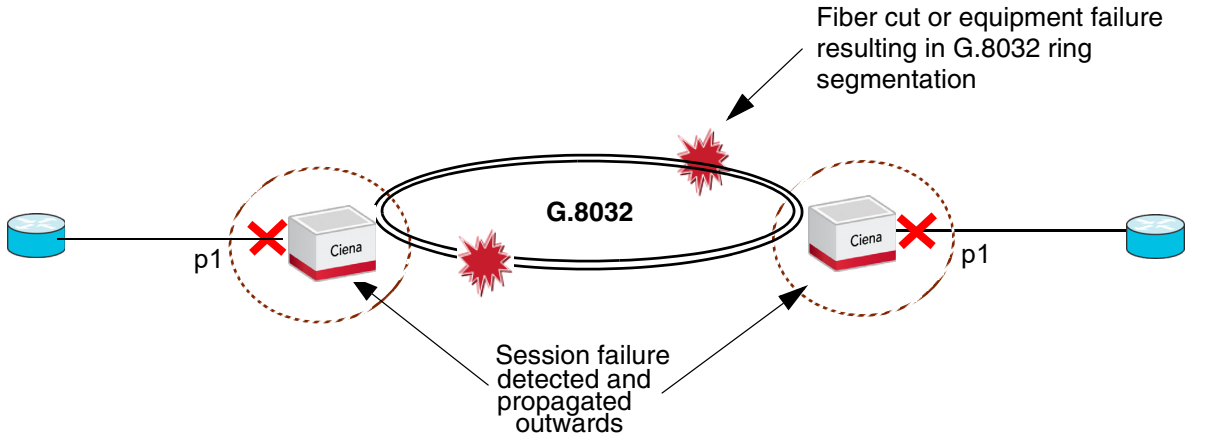

This figure shows a fiber cut or equipment failure that impacts the connection between a 39XX/51XX switch and a CPE through a G.8032 ring.

Copyright© 2022 Ciena® Corporation

Figure 19-9
Connection between 39XX/51XX switch and CPE through a G.8032 ring lost due to fiber cut or equipment failure

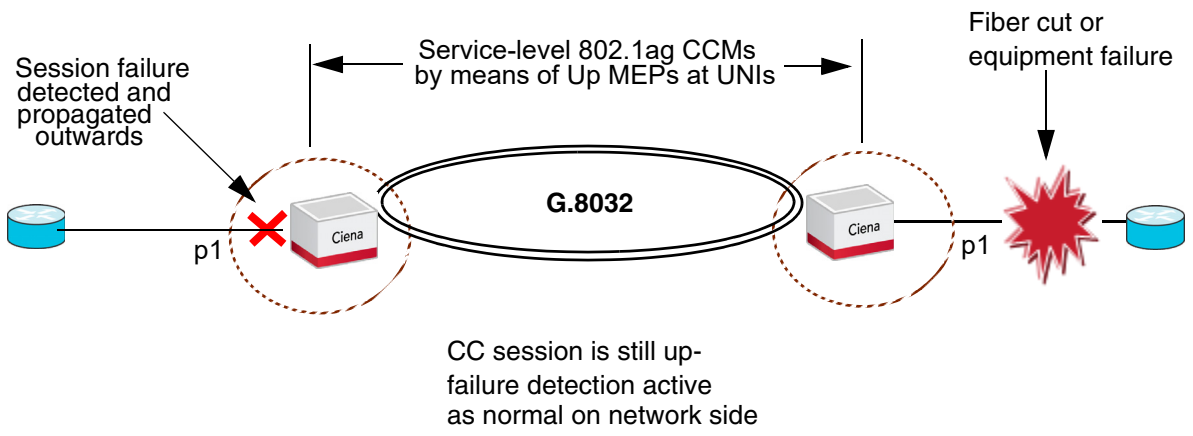

Step	Action
1	Enable CFM: `cfm enable`

2	Create a CFM service:
	`cfm service create vs <vs> name <String[45]> level <NUMBER: 0-7> Enable the CFM service created in step 2:`
	`cfm service enable service <service>`

3	Create a MEP for the CFM service: `cfm mep create service <service> sub-port <sub-port> type up mepid <NUMBER: 1-8191>`
4	Assign an alarm priority to the CFM service: `cfm service set service <service> alarm-priority <NUMBER: 1-5>`
5	Create a VLLI instance group: `virtual-link-loss-indication create instance-group <String[31]>`
6	Create a VLLI action group: `virtual-link-loss-indication create action-group <String[15]>`
7	Modify the instance group created in step 5 to set the source mode action to the CFM service: `virtual-link-loss-indication add instance-group <instance-group> mode source cfm-instance <service>`

Copyright© 2022 Ciena® Corporation

8 Modify the instance group to set the destination mode port instance:

```
virtual-link-loss-indication add instance-group
<instance-group> mode destination port-instance <Port
list>
```

9 Modify the action group to set the fault action to port-shut:

```
virtual-link-loss-indication add action-group <action-
group> type fault action port-shut precedence <NUMBER: 1-
8>
```

10 Modify the action group to set the recovery action to port-unshut:

```
virtual-link-loss-indication add action-group <action-
group> type recovery action port-unshut precedence
<NUMBER: 1-8>
```

11 Modify the instance group to set the direction:

```
virtual-link-loss-indication set instance-group
<instance-group> direction unidirectional
```

12 Modify the instance group to set the port instance and assign the action-group:

```
virtual-link-loss-indication set instance-group
<instance-group> port-instance <Port list> action-group
<action-group>
```

13 Enable the instance group:

```
virtual-link-loss-indication enable instance-group
<instance-group>
```

—end—

Copyright© 2022 Ciena® Corporation

Example

This example creates

- a CFM service CFM_Svc_1 on virtual switch VS_1
- instance group g1
- action group a1

As a result, if the service CFM_Svc_1 detects CCM loss, VLLI triggers the associated action, which is to shut the port on a fault. VLLI restores the port on recovery.

Refer to the Connection between 39XX/51XX switches lost due to fiber cut or equipment failure figure for configuration details.

```
cfm enable
cfm service create vs VS_1 name CFM_Svc_1 level 4
cfm service enable service CFM_Svc_1
cfm mep create service CFM_Svc_1 port 1 type up mepid 200
cfm service set service CFM_Svc_1 alarm-priority 2

virtual-link-loss-indication create instance-group g1
virtual-link-loss-indication create action-group a1

virtual-link-loss-indication add instance-group g1 mode source
cfm-instance CFM_Svc_1
virtual-link-loss-indication add instance-group g1 mode
destination port-instance 1

virtual-link-loss-indication add action-group a1 type fault
action port-shut precedence 1
virtual-link-loss-indication add action-group a1 type recovery
action port-unshut precedence 1

virtual-link-loss-indication set instance-group g1 direction
unidirectional
virtual-link-loss-indication set instance-group g1 port-
instance 1 action-group a1

virtual-link-loss-indication enable instance-group g1
```


Copyright© 2022 Ciena® Corporation

Procedure 19-2
Configuring VLLI to shut down a port when CCM loss is detected - VLAN

Configure VLLI to automatically shut down a port when CCM loss is detected. This figure shows a sample scenario for shutting down a port when VLLI detects CCM loss. This scenario is based on a CFM service running on a VLAN.

Figure 19-10
Connection between 39XX/51XX switches lost due to fiber cut or equipment failure

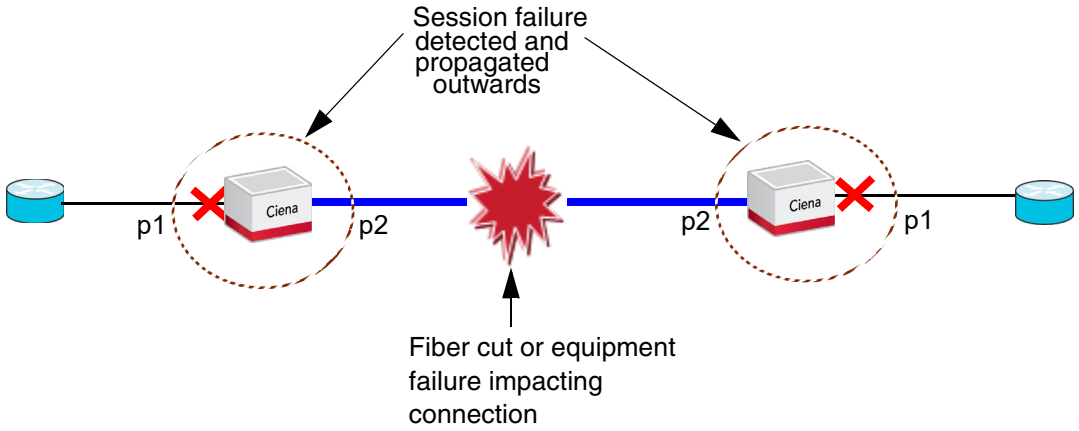

Step	Action
1	Enable CFM: `cfm enable`

2	Create a CFM service:
	`cfm service create vlan <vlan> name <String[45]> level <NUMBER: 0-7> Enable the CFM service created in step 2:`
	`cfm service enable service <service>`

3	Create a MEP for the CFM service: `cfm mep create service <service> sub-port <sub-port> mepid <NUMBER: 1-8191> type down`
4	Create a VLLI instance group: `virtual-link-loss-indication create instance-group <String[31]>`
5	Create a VLLI action group (multiple action groups can be created): `virtual-link-loss-indication create action-group <String[15]>`

Copyright© 2022 Ciena® Corporation

6 Modify the instance group created in step 4 to set the source mode action to the CFM service:

```
virtual-link-loss-indication add instance-group
<instance-group> mode source cfm-instance <service>
```

7 Modify the instance group to set the destination mode port instance:

```
virtual-link-loss-indication add instance-group
<instance-group> mode destination port-instance <Port
list>
```

8 Modify the action group to set a fault action to port-shut:

```
virtual-link-loss-indication add action-group <action-
group> type fault action port-shut precedence <NUMBER: 1-
8>
```

9 Modify the action group to set the recovery action to port-unshut:

```
virtual-link-loss-indication add action-group <action-
group> type recovery action port-unshut precedence
<NUMBER: 1-8>
```

10 Modify the action group to set a fault action to ccm-stop:

```
virtual-link-loss-indication add action-group <action-
group> type fault action ccm-stop precedence <NUMBER: 1-
8>
```

11 Modify the action group to set the recovery action to ccm-resume:

```
virtual-link-loss-indication add action-group <action-
group> type recovery action ccm-resume precedence
<NUMBER: 1-8>
```

12 Modify the instance group to set the direction:

```
virtual-link-loss-indication set instance-group
<instance-group> direction bidirectional
```

13 Modify the instance group to set the cfm instance and assign the action-group:

```
virtual-link-loss-indication set instance-group
<instance-group> cfm-instance <Port list> action-group
<action-group>
```

14 Modify the instance group to set the port instance and assign the action-group:

```
virtual-link-loss-indication set instance-group
<instance-group> port-instance <Port list> action-group
<action-group>
```

15 Enable the instance group:

```
virtual-link-loss-indication enable instance-group
<instance-group>
```

—end—

Copyright© 2022 Ciena® Corporation

Example

This example creates

- a CFM service CFM_Svc_VLAN100 on virtual LAN (VLAN) 100
- instance group g1
- action groups a1 and a2

As a result, if the service CFM_Svc_VLAN100 detects CCM loss, VLLI triggers the associated actions, which is to shut the port on a fault and stop CCM transmissions. On recovery, VLLI restores the port and restarts CCM transmissions.

```
cfm enable
cfm service create vlan 100 name CFM_Svc_VLAN100 level 3
cfm service enable service CFM_Svc_VLAN100
cfm mep create service CFM_Svc_VLAN100 port 2 vlan 100 mepid 200
type down

virtual-link-loss-indication create instance-group g1
virtual-link-loss-indication create action-group a1
virtual-link-loss-indication create action-group a2

virtual-link-loss-indication add instance-group g1 mode source
cfm-instance CFM_Svc_VLAN100
virtual-link-loss-indication add instance-group g1 mode
destination port-instance 1

virtual-link-loss-indication add action-group a1 type fault
action port-shut precedence 1
virtual-link-loss-indication add action-group a1 type recovery
action port-unshut precedence 1

virtual-link-loss-indication add action-group a2 type recovery
action ccm-resume precedence 1
virtual-link-loss-indication add action-group a2 type fault
action ccm-stop precedence 1

virtual-link-loss-indication set instance-group g1 direction
bidirectional
virtual-link-loss-indication set instance-group g1 cfm-instance
CFM_Svc_VLAN100 action a2
virtual-link-loss-indication set instance-group g1 port-
instance 1 action a1

virtual-link-loss-indication enable instance-group g1
```


Copyright© 2022 Ciena® Corporation

Procedure 19-3
Configuring VLLI between a UNI and an ENNI

Configure VLLI to automatically shut down a port when the connection between a UNI and an ENNI is impacted.

This figure shows an example of the connection between a UNI and an ENNI being impacted.

Figure 19-11
Loss of connection between UNI and ENNI

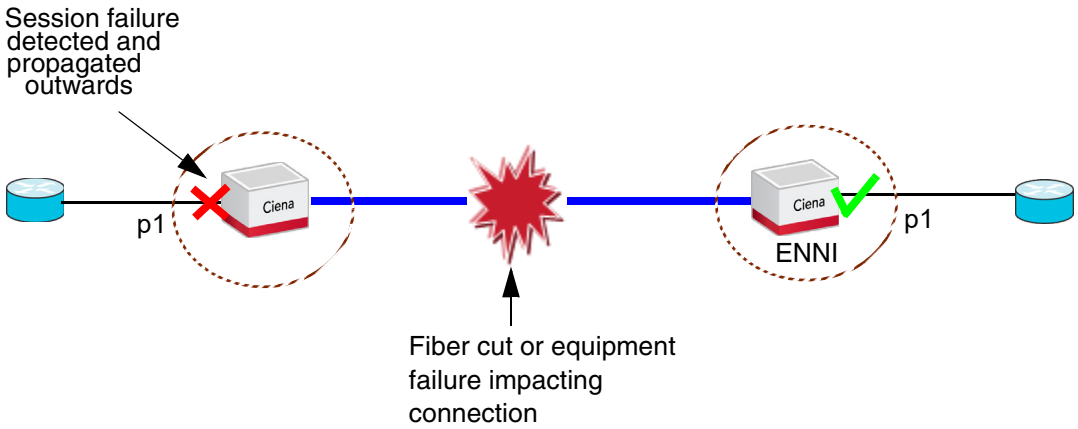

Step	Action

To configure CFM service at UNI

1 Enable CFM:

```
cfm enable
```

2 Create a CFM service:

```
cfm service create vs <vs> name <String[45]> level
<NUMBER: 0-7>
```

3 Enable the CFM service created in step 2:

```
cfm service enable service <service>
```

4 Create a MEP for the CFM service:

```
cfm mep create service <service> sub-port <sub-port> type
up mepid <NUMBER: 1-8191>
```

5 Assign an alarm priority to the CFM service:

```
cfm service set service <service> alarm-priority <NUMBER:
1-5>
```


Copyright© 2022 Ciena® Corporation

To configure CFM service at ENNI

6 Enable CFM:

```
cfm enable
```

7 Create a CFM service:

```
cfm service create vs <vs> name <String[45]> level
<NUMBER: 0-7>
```

8 Enable the CFM service created in step 2:

```
cfm service enable service <service>
```

9 Create a MEP for the CFM service:

```
cfm mep create service <service> sub-port <sub-port> type
up mepid <NUMBER: 1-8191>
```

To configure VLLI

10 Create a VLLI instance group:

```
virtual-link-loss-indication create instance-group
<String[31]>
```

11 Create a VLLI action group:

```
virtual-link-loss-indication create action-group
<String[15]>
```

12 Modify the instance group created in step 10 to set the source mode action to the CFM service:

```
virtual-link-loss-indication add instance-group
<instance-group> mode source cfm-instance <service>
```

13 Modify the instance group to set the destination mode port instance:

```
virtual-link-loss-indication add instance-group
<instance-group> mode destination port-instance <Port
list>
```

14 Modify the action group created in step 11 to set the fault action to port-shut:

```
virtual-link-loss-indication add action-group <action-
group> type fault action port-shut precedence <NUMBER: 1-
8>
```

15 Modify the action group to set the recovery action to port-unshut:

```
virtual-link-loss-indication add action-group <action-
group> type recovery action port-unshut precedence
<NUMBER: 1-8>
```

16 Modify the instance group to set the direction:

```
virtual-link-loss-indication set instance-group
<instance-group> direction unidirectional
```


Copyright© 2022 Ciena® Corporation

17 Modify the instance group to set the port instance and assign the action-group:

```
virtual-link-loss-indication set instance-group
<instance-group> port-instance <Port list> action-group
<action-group>
```

18 Enable the instance group:

```
virtual-link-loss-indication enable instance-group
<instance-group>
```

—end—

Example

This example creates

- a CFM service CFM_Svc_1 on virtual switch VS_1
- instance group g1
- action group a1

As a result, if the service CFM_Svc_1 detects CCM loss, VLLI triggers the associated action, which is to shut the port on a fault. VLLI unshuts the port on recovery.

This example creates the CFM service on the UNI.

```
cfm enable
cfm service create vs VS_1 name CFM_Svc_1 level 4
cfm service enable service CFM_Svc_1
cfm mep create service CFM_Svc_1 port 1 type up mepid 200
cfm service set service CFM_Svc_1 alarm-priority 2
```

This example creates the CFM service on the ENNI.

```
cfm enable
cfm service create vs VS_1 name CFM_Svc_1 level 4
cfm service enable service CFM_Svc_1
cfm mep create service CFM_Svc_1 port 1 type up mepid 200
```


Copyright© 2022 Ciena® Corporation

This example configures VLLI.

```
virtual-link-loss-indication create instance-group g1
virtual-link-loss-indication create action-group a1

virtual-link-loss-indication add instance-group g1 mode source
cfm-instance CFM_Svc_1
virtual-link-loss-indication add instance-group g1 mode
destination port-instance 1

virtual-link-loss-indication add action-group a1 type fault
action port-shut precedence 1
virtual-link-loss-indication add action-group a1 type recovery
action port-unshut precedence 1

virtual-link-loss-indication set instance-group g1 direction
unidirectional
virtual-link-loss-indication set instance-group g1 port-
instance 1 action-group a1

virtual-link-loss-indication enable instance-group g1
```


Copyright© 2022 Ciena® Corporation

Procedure 19-4
Configuring inverse VLLI with one 39XX/51XX switch and two PEs

In an inverse VLLI configuration, one link is active and one link is backup. If the active link fails, the backup link becomes active.

When two CPEs are connected to one 39XX/51XX switch in an inverse VLLI configuration, there is one active link, p1, and one backup link, p2. The 39XX/51XX switch maintains the p2 link in the down state from the perspective of CPE2. If the p1 link fails, the 39XX/51XX switch executes a failover, which causes the p2 link to come up.

If the p1 link is restored, the 39XX/51XX switch fails back, which forces the p1 link to change status to up while shutting down the p2 link.

Step	Action
1	Create an instance group: `virtual-link-loss-indication create instance-group <String[31]>`
2	Create an action group: `virtual-link-loss-indication create action-group <String[15]>`
3	Modify the instance group created in step 1 to set the source mode action to the port instance: `virtual-link-loss-indication add instance-group <instance-group> mode source port-instance <Port list>`
4	Modify the instance group to set the source mode action to the port instance: `virtual-link-loss-indication add instance-group <instance-group> mode destination port-instance <Port list>`
5	Modify the action group created in step 2 to set the fault action to port-unshut: `virtual-link-loss-indication add action-group <action-group> type fault action port-unshut`
6	Modify the action group to set the recovery action to port-shut: `virtual-link-loss-indication add action-group <action-group> type recovery action port-shut`
7	Modify the instance group to set the direction to unidirectional: `virtual-link-loss-indication set instance-group <instance-group> direction unidirectional`

Copyright© 2022 Ciena® Corporation

8 Modify the instance group to set the port instance:

```
virtual-link-loss-indication set instance-group
<instance-group> port-instance <Port list> action
<action-group>
```

9 Enable the instance group:

```
virtual-link-loss-indication enable instance-group
<instance-group>
```

—end—

Example

In this scenario, the PEs are connected to one 39XX/51XX switch with two links: link p1, which is the active link, and link p2, which is the backup link. The 39XX/51XX switch maintains link p2 in the down state from the perspective of PE2. If link p1 fails, the 39XX/51XX switch executes a failover, which causes link p2 to come up. If link p1 is restored, the 39XX/51XX switch fails back, which forces PE1 link to change status to up while shutting down the PE2 connection.

This figure illustrates the connections.

Figure 19-12
Inverse VLLI example

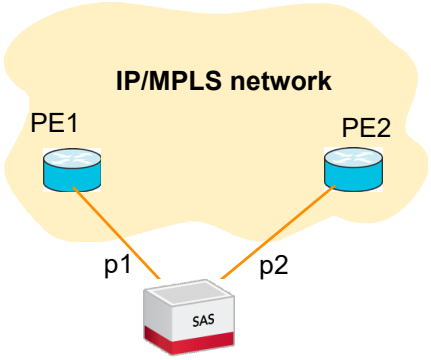

Copyright© 2022 Ciena® Corporation

This example configures inverse VLLI.

```
virtual-link-loss-indication create instance-group g1
virtual-link-loss-indication create action-group a1

virtual-link-loss-indication add instance-group g1 mode source
port-instance 1
virtual-link-loss-indication add instance-group g1 mode
destination port-instance 2

virtual-link-loss-indication add action-group a1 type fault
action port-unshut
virtual-link-loss-indication add action-group a1 type recovery
action port-shut

virtual-link-loss-indication set instance-group g1 direction
unidirectional
virtual-link-loss-indication set instance-group g1 port-
instance 2 action a1

virtual-link-loss-indication enable instance-group g1
```


Copyright© 2022 Ciena® Corporation

Procedure 19-5
Configuring inverse VLLI with two 39XX/51XX switches and two PEs

In an inverse VLLI configuration, one link is active and one link is backup. If the active link fails, the backup link becomes active.

When two CPEs are connected to two 39XX/51XX switches in an inverse VLLI configuration, there is one active link, p1, and one backup link, p2, for each 39XX/51XX switch. The 39XX/51XX switch maintains the p2 link in the down state from the perspective of the CPE. If the p1 link fails, the 39XX/51XX switch executes a failover, which causes the p2 link to come up.

If the p1 link is restored, the 39XX/51XX switch fails back, which forces the p1 link to change status to up while shutting down the p2 link.

Step	Action
1	Enable CFM: `cfm enable`
2	Create a CFM service: `cfm service create vs <vs> name <String[45]> level <NUMBER: 0-7>`
3	Enable the CFM service created in step 2: `cfm service enable service <service>`
4	Create a MEP for the CFM service: `cfm mep create service <service> sub-port <sub-port> type up mepid <ID1>`
5	Assign an alarm priority to the CFM service: `cfm service set service CFM_Svc_1 alarm-priority <NUMBER: 1-5>`
6	Create an instance group: `virtual-link-loss-indication create instance-group <String[31]>`
7	Create an action group: `virtual-link-loss-indication create action-group <String[15]>`
8	Modify the instance group to set the source mode action to the CFM service instance: `virtual-link-loss-indication add instance-group <instance-group> mode source cfm-instance <service>`

Copyright© 2022 Ciena® Corporation

9 Modify the instance group to set the destination mode port instance:

```
virtual-link-loss-indication add instance-group
<instance-group> mode destination port-instance <Port
list>
```

10 Modify the action group to set the fault action to port-unshut and to set the precedence:

```
virtual-link-loss-indication add action-group <action-
group> type fault action port-unshut precedence <NUMBER:
1-8>
```

11 Modify the action group to set the recovery action to port-shut and to set the precedence:

```
virtual-link-loss-indication add action-group <action-
group> type recovery action port-shut precedence <NUMBER:
1-8>
```

12 Modify the instance group to set the direction to unidirectional:

```
virtual-link-loss-indication set instance-group
<instance-group> direction unidirectional
```

13 Modify the instance group to set the port instance and assign the action-group:

```
virtual-link-loss-indication set instance-group
<instance-group> port-instance <Port list> action-group
<action-group>
```

14 Enable the instance-group:

```
virtual-link-loss-indication enable instance-group
<instance-group>
```

—end—

Copyright© 2022 Ciena® Corporation

Example

This figure shows an inverse VLLI example with two 39XX/51XX switches and two CPEs.

Figure 19-13
Inverse VLLI example with two 39XX/51XX switches and two CPEs

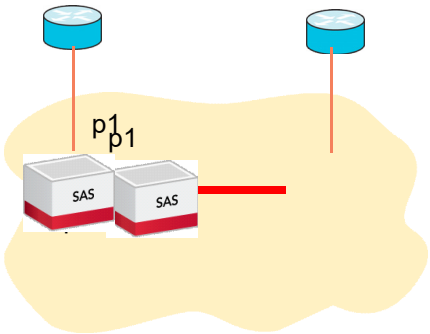

This example configures the example shown in Figure 19-13.

```
cfm enable
cfm service create vs VS_1 name CFM_Svc_1 level 4
cfm service enable service CFM_Svc_1
cfm mep create service CFM_Svc_1 port 1 type up mepid 200
cfm service set service CFM_Svc_1 alarm-priority 2

virtual-link-loss-indication create instance-group g1
virtual-link-loss-indication create action-group a1

virtual-link-loss-indication add instance-group g1 mode source
cfm-instance CFM_Svc_1
virtual-link-loss-indication add instance-group g1 mode
destination port-instance 1

virtual-link-loss-indication add action-group a1 type fault
action port-unshut precedence 1
virtual-link-loss-indication add action-group a1 type recovery
action port-shut precedence 1

virtual-link-loss-indication set instance-group g1 direction
unidirectional
virtual-link-loss-indication set instance-group g1 port-
instance 1 action-group a1

virtual-link-loss-indication enable instance-group g1
```


Copyright© 2022 Ciena® Corporation

Procedure 19-6
Configuring VLLI for MAC flush on an intersecting node

In this scenario, intersecting node D does not have ring protection and it is unidirectional.

Figure 19-14
VLLI configured on intersecting node D for MAC flush

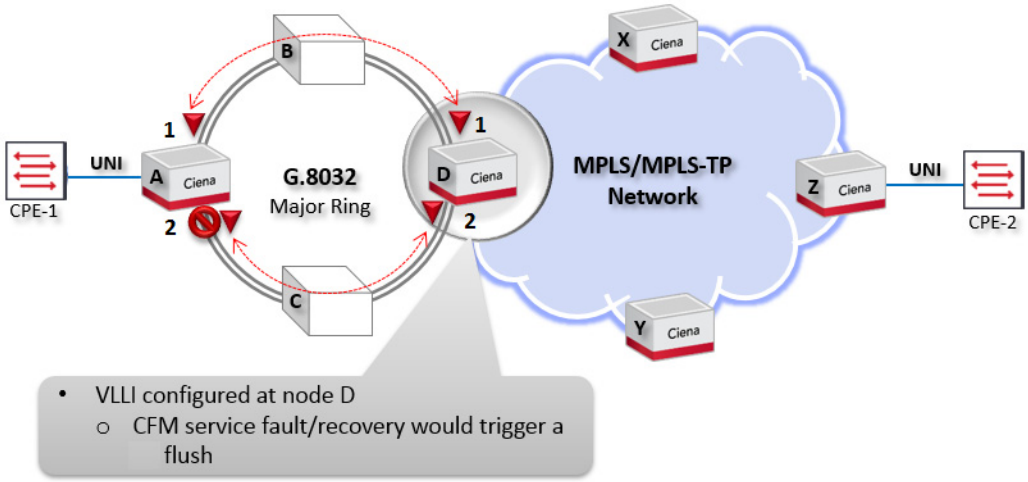

Step	Action
1	Protect the upper half ring (nodes A,B,D) with a protected VLAN and the lower half of the ring (nodes A,C,D) with another protected VLAN.
2	Create a CFM service on those protected VLANs. And configure a down MEPS between the RPL owner and the intersecting node (node D).
3	As the port 2 of node A is the RPL owner, at first the CFM service of the lower half will be in fault and the upper half will be in an OK state.
4	As soon as a fault occurs in the upper half of the ring, port 1 of node D is flushed due to mac-flush action and the traffic moves to the lower half of the ring where the RPL is removed and the CFM service covering the lower half of the ring moves to a recovery state.
5	As soon as the upper half of the ring recovers, the lower half again goes into fault due to RPL imposition, hence, flushing port 2 of node D.

—end—

Copyright© 2022 Ciena® Corporation

Example

This example configures the intersecting node D shown in Figure 19-14. The R-APS-protected VID for this example are 100, 300.

> ***Note:*** This example is just for the intersecting node D, which is a small part of a larger G.8032 configuration.

```
cfm enable
cfm service create vlan 100 name S1
cfm service enable service S1
cfm mep create service S1 port 1 type down mepid 100
cfm service set service S1 alarm-priority 2 alarm-time 0

virtual-link-loss-indication create instance-group g1
virtual-link-loss-indication create action-group a1

virtual-link-loss-indication add instance-group g1 mode source
cfm-instance S1
virtual-link-loss-indication add instance-group g1 mode
destination port-instance 1

virtual-link-loss-indication add action-group a1 type fault
action mac-flush precedence 1
virtual-link-loss-indication add action-group a1 type recovery
action mac-flush precedence 1

cfm enable
cfm service create vlan 300 name S2
cfm service enable service S2
cfm mep create service S2 port 2 type down mepid 300
cfm service set service S2 alarm-priority 2 alarm-time 0

virtual-link-loss-indication create instance-group g2
virtual-link-loss-indication create action-group a2

virtual-link-loss-indication add instance-group g2 mode source
cfm-instance S2
virtual-link-loss-indication add instance-group g2 mode
destination port-instance 2

virtual-link-loss-indication add action-group a2 type fault
action mac-flush precedence 1
virtual-link-loss-indication add action-group a2 type recovery
action mac-flush precedence 1
```


Copyright© 2022 Ciena® Corporation

Procedure 19-7
Configuring MAC flush from the MPLS network ingress node to all other nodes

This figure depicts a scenario where MAC flush is configured from the MPLS network ingress node to all other nodes:

Figure 19-15
MAC flush from the MPLS network ingress node to all other nodes

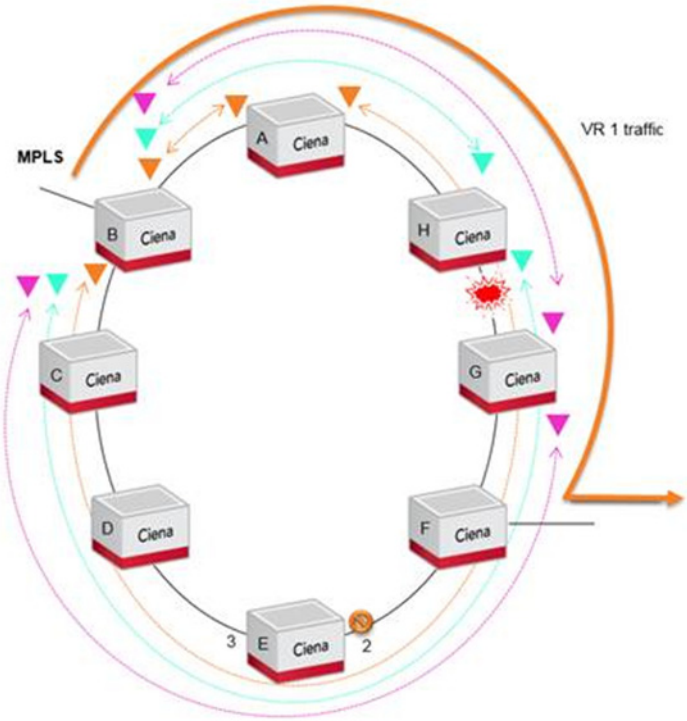

In this scenario, the following has occurred:

- VR 1 has node E as the RPL owner. The CFM service protecting the half ring (nodes B,A,H,G,F,E) is in MacStatus/PortStatus fault.
- Unidirectional traffic flows from node B to node F.
- The link between nodes H and G breaks.
- The CFM service monitoring this link does not register any fault because the service was already in a faulty state.
- Node B never flushes its ports, thus leading to a black hole for traffic.

Copyright© 2022 Ciena® Corporation

Recommended solution

Configure a CFM service from the MPLS network ingress node B to all other nodes.

Now if the link breaks between nodes H and G, the green and pink CFM services register it and node B can flush its ports.

Copyright© 2022 Ciena® Corporation

Procedure 19-8
Y-cable client protection using inverse VLLI between nodes B and A

Use this procedure to configure inverse VLLI at a y-protected node.

The following figure depicts a topology with Y-cable client protection between two nodes, identified as node A and node B.

Figure 19-16
Y-cable client protection between nodes B and A

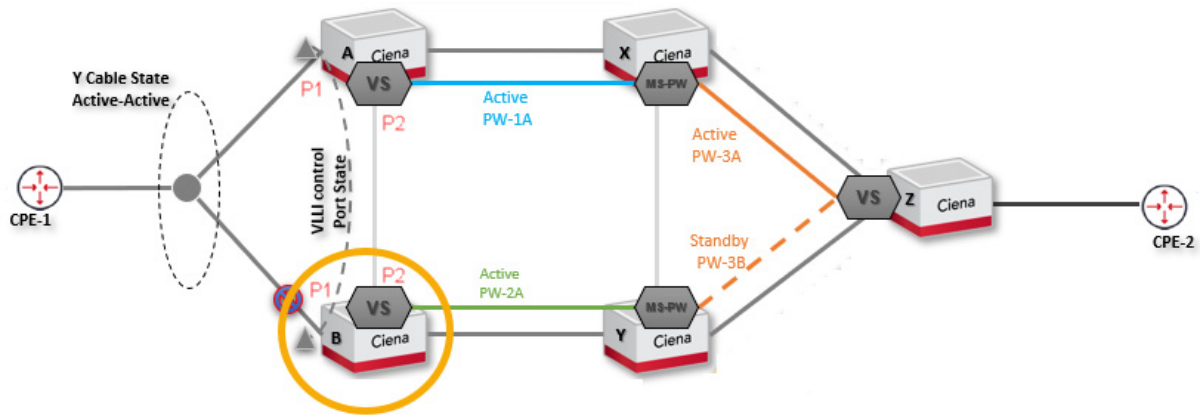

The objective is to use inverse VLLI to create Y-cable client protection between two y-protected nodes, identified in the figure as node A and node B.

All commands are entered at the node where inverse VLLI is being configured--node B in the example.

When the CFM fault is received, action port-unshut is applied to the port for the split end of the Y-cable at node B. When the fault is cleared, action port-shut is applied port for the split end of the Y-cable at node B.

Prerequisites

A sub-port has been created and has been attached to a virtual switch. You know the names of the sub-port and the virtual switch, and you know the name of the sub-port's parent port.

Copyright© 2022 Ciena® Corporation

Step	Action

Create the CFM service

1 Set the y-protection role of the port.

```
port set port <port-name> y-protection
<primary|secondary>
```

2 Enable CFM.

```
cfm enable
```

3 Create the CFM service.

```
cfm service create vs <vs-name> name <cfm-service> level
<lvl>
```

4 Enable the CFM service.

```
cfm service enable service <cfm-service>
```

5 For the CFM service, set the alarm priority, set CFM alarm time, set CFM service reset time, and enable y-protection monitoring.

```
cfm service set service <cfm-service> alarm-priority 2
alarm-time 0 reset-time 100 y-protection-monitoring on
```

Note 1: CFM service alarm time is the duration that CFM waits before notifying external clients (such as VLLI, G.8032, alarm engine). For Y-cable, the recommended alarm-time value is 0 milliseconds. The value should be configured in increments of 100 milliseconds. The default is three seconds.

Note 2: CFM service reset time is the duration between the service fault clear event and the time that CFM notifies external clients (such as VLLI, G.8032, alarm engine). For Y-cable, the recommended reset time is 100 milliseconds. The default is three seconds. If the service fault occurs again during the reset time, there is no notification to external clients.

6 Create the mep.

```
cfm mep create service <cfm-service> sub-port <sub-port-
name> type up mepid <mepid>
```

where

port-name is the name of the y-protected port

Note: As shown in this procedure, you must set the y-protection monitoring parameter of a CFM service before creating the MEP. The y-protection monitoring parameter cannot be modified if a MEP exists for the service. If you want to modify the y-protection monitoring parameter of a CFM service, you must first delete all MEPs that have been created for that service.

Copyright© 2022 Ciena® Corporation

Configure inverse VLLI

7 Create an instance group.

```
virtual-link-loss-indication create instance-group g1
```

8 Create an action group.

```
virtual-link-loss-indication create action-group a1
```

9 Specify that the CFM service is the source for the instance group.

```
virtual-link-loss-indication add instance-group g1 mode
source cfm-instance <cfm-service>
```

10 Specify that the y-protected port on the node is the destination for the instance group.

```
virtual-link-loss-indication add instance-group g1 mode
destination port-instance <port-name>
```

11 Specify the fault and recovery actions for the action group.

```
virtual-link-loss-indication add action-group a1 type
fault action port-unshut precedence 1
```

```
virtual-link-loss-indication add action-group a1 type
recovery action port-shut precedence 1
```

12 Specify that the direction for the instance group is unidirectional.

```
virtual-link-loss-indication set instance-group g1
direction unidirectional
```

13 Specify that the actions for the action group apply to the y-protected port and belong to the instance group.

```
virtual-link-loss-indication set instance-group g1 port-
instance <port-name> action a1
```

14 Enable the instance group.

```
virtual-link-loss-indication enable instance-group g1
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 19-9
MPLS active/active protection state management using VLLI in a Y-cable topology

Use this procedure to configure MPLS active/active protection state management can using VLLI in a Y-cable topology.

The following figure shows the topology.

Figure 19-17
Y-cable topology for MPLS active/active protection state management using VLLI

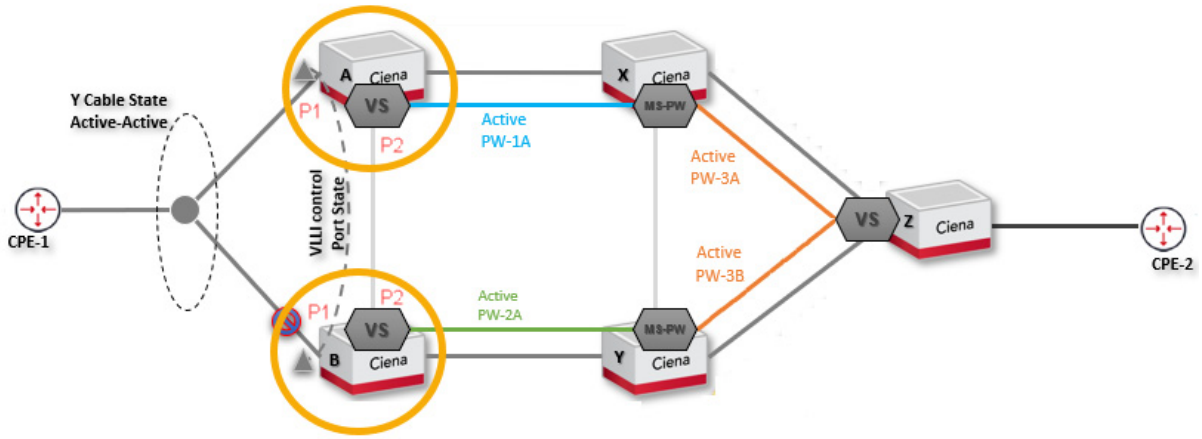

In the active/active model, each split end of the Y-cable connects to an active pseudo wire (PW). In the example, the Y-cable at the left connects to nodes A and B. At each of these nodes, the split end connects to an active PW: PW-1A at node A and PW-2A at node B.

In this topology, the primary node is the node with the split-end port configured as y-protection primary. The secondary node is the node with the split-end port configured as y-protection secondary. VLLI decides which split-end port will be used. After startup or after a restart, VLLI selects the y-protection primary port. Subsequently VLLI changes ports if a fault requires it to do so, and auto-reversion does not occur.

Prerequisites

At each node, primary and secondary, a sub-port has been created and has been attached to a virtual switch. You know the names of the sub-ports and the virtual switches, and you know the names of the parent ports, and their port numbers (to be used as the values of the “port-instance” parameter).

Copyright© 2022 Ciena® Corporation

Step	Action
1	At the primary node, specify that the port connected to the split end of the Y-cable is the primary port. `port set port <port-name> y-protection primary`
2	At the primary node, enable CFM. `cfm enable`
3	At the primary node, create the CFM service on a virtual switch that uses the y-protected port. `cfm service create vs <vs-name> name <cfm-service> level <lvl>`
4	At the primary node, enable the CFM service. `cfm service enable service <cfm-service>`
5	For the CFM service, set the alarm priority, set the CFM alarm time, set the CFM service reset time, and enable y-protection monitoring. `cfm service set service <cfm-service> alarm-priority 2 alarm-time 0 reset-time 100 y-protection-monitoring on` ***Note 1:*** CFM service alarm time is the duration (in milliseconds) that CFM waits before sending a fault alarm to external clients (such as VLLI, G.8032, alarm engine). For Y-cable, the recommended alarm-time value is 0 milliseconds. The value should be configured in increments of 100 milliseconds. The default is three seconds. ***Note 2:*** CFM service reset time is the duration between the service fault clear event and the time that CFM notifies external clients (such as VLLI, G.8032, alarm engine). For Y-cable, the recommended reset time is 100 milliseconds. The default is three seconds. The default is three seconds. If the service fault occurs again during the reset time, there is no notification to external clients.
6	At the primary node, create the mep on the sub-port. `cfm mep create service <cfm-service> sub-port <sub-port-name> type up mepid <mepid>` where port-name is the name of the y-protected port ***Note:*** As shown in this procedure, you must set the y-protection monitoring parameter of a CFM service before creating the MEP. The y-protection monitoring parameter cannot be modified if a MEP exists for the service. If you want to modify the y-protection monitoring parameter of a CFM service, you must first delete all MEPs that have been created for that service.
7	At the primary node, create an action group (the first action group). `virtual-link-loss-indication create action-group <first action group>`

Copyright© 2022 Ciena® Corporation

8 At the primary node, specify that the fault action for the first action group is port-unshut.

```
virtual-link-loss-indication add action-group <first
action group> type fault action port-unshut precedence 1
```

9 At the primary node, specify that the recovery action for the first action group is port-shut.

```
virtual-link-loss-indication add action-group <first
action group> type recovery action port-shut precedence 1
```

10 At the primary node, create a second action group.

```
virtual-link-loss-indication create action-group <second
action group>
```

11 At the primary node, specify that the fault action for the second action group is port-shut.

```
virtual-link-loss-indication add action-group <second
action group> type fault action port-shut precedence 1
```

12 At the primary node, specify that the recovery action for the second action group is port-unshut.

```
virtual-link-loss-indication add action-group <second
action group> type recovery action port-unshut precedence
1
```

13 At the primary node, create a third action group.

```
virtual-link-loss-indication create action-group <third
action group>
```

14 At the primary node, specify that the recovery action for the third action group is vc-mac-withdraw.

```
virtual-link-loss-indication add action-group <third
action group> type recovery action vc-mac-withdraw
precedence 1
```

Note: The vc-mac-withdraw message trigger by inverse VLLI (iVLLI) is supported only for static PWs.

15 At the primary node, create an instance group (the first instance group)

```
virtual-link-loss-indication create instance-group
<first instance group>
```

16 At the primary node, specify that the CFM service is the source for the first instance group.

```
virtual-link-loss-indication add instance-group <first
instance group> cfm-instance <cfm-service> mode source
```


Copyright© 2022 Ciena® Corporation

17 At the primary node, specify that the port-instance is the destination for the first instance group, and that the actions on the port-instance are those specified by the first action group. (The first action group specifies fault and recovery actions.)

```
virtual-link-loss-indication add instance-group <first
instance group> port-instance <port number> mode
destination action-group <first action group>
```

18 At the primary node, enable the first instance group.

```
virtual-link-loss-indication enable instance-group
<first instance group>
```

19 At the primary node, create a second instance group, specifying that it is bidirectional.

```
virtual-link-loss-indication create instance-group
<second instance group> direction bidirectional
```

20 At the primary node, specify that the vc-instance is the source for the second instance group, and that the actions on the vc-instance are those specified by the third action group. (The third action group specifies a recovery action.)

```
virtual-link-loss-indication add instance-group <second
instance group> vc-instance <vc instance> mode source
action-group <third action group>
```

Note: Actions are applicable to the source because the instance group is bidirectional.

21 At the primary node, specify that the port-instance is the destination for the second instance group, and that the actions on the port-instance are those specified by the second action group. (The second action group specifies fault and recovery actions.)

```
virtual-link-loss-indication add instance-group <second
instance group> port-instance <port number> mode
destination action-group <second action group>
```

22 At the primary node, enable the second instance group.

```
virtual-link-loss-indication enable instance-group
<second instance group>
```

23 At the secondary node, specify that the port connected to the split end of the Y-cable is the secondary port.

```
port set port <port-name> y-protection secondary
```

24 At the secondary node, enable CFM.

```
cfm enable
```

25 At the secondary node, create the CFM service on a virtual switch that uses the y-protected port.

```
cfm service create vs <vs-name> name <cfm-service> level
<lvl>
```

26 At the secondary node, enable the CFM service.

```
cfm service enable service <cfm-service>
```


Copyright© 2022 Ciena® Corporation

27 For the CFM service, set the alarm priority, set the CFM alarm time, set the CFM service reset time, and enable y-protection monitoring.

```
cfm service set service <cfm-service> alarm-priority 2
alarm-time 0 reset-time 100 y-protection-monitoring on
```

Note 1: CFM service alarm time is the duration (in milliseconds) that CFM waits before sending a fault alarm to external clients (such as VLLI, G.8032, alarm engine). For Y-cable, the recommended alarm-time value is 0 milliseconds. The value should be configured in increments of 100 milliseconds. The default is three seconds.

Note 2: CFM service reset time is the duration between the service fault clear event and the time that CFM notifies external clients (such as VLLI, G.8032, alarm engine). For Y-cable, the recommended reset time is 100 milliseconds. The default is three seconds. The default is three seconds. If the service fault occurs again during the reset time, there is no notification to external clients.

28 At the secondary node, create the mep on the sub-port.

```
cfm mep create service <cfm-service> sub-port
<sub-port-name> type up mepid <mepid>
```

where

port-name is the name of the y-protected port

Note: As shown in this procedure, you must set the y-protection monitoring parameter of a CFM service before creating the MEP. The y-protection monitoring parameter cannot be modified if a MEP exists for the service. If you want to modify the y-protection monitoring parameter of a CFM service, you must first delete all MEPs that have been created for that service.

29 At the secondary node, create an action group (the first action group).

```
virtual-link-loss-indication create action-group <first
action group>
```

30 At the secondary node, specify that the fault action for the first action group is port-unshut.

```
virtual-link-loss-indication add action-group <first
action group> type fault action port-unshut precedence 1
```

31 At the secondary node, specify that the recovery action for the first action group is port-shut.

```
virtual-link-loss-indication add action-group <first
action group> type recovery action port-shut precedence 1
```

32 At the secondaryy node, create a second action group.

```
virtual-link-loss-indication create action-group <second
action group>
```

33 At the secondary node, specify that the fault action for the second action group is port-shut.

```
virtual-link-loss-indication add action-group <second
action group> type fault action port-shut precedence 1
```


Copyright© 2022 Ciena® Corporation

34 At the secondary node, specify that the recovery action for the second action group is port-unshut.

```
virtual-link-loss-indication add action-group <second
action group> type recovery action port-unshut precedence
1
```

35 At the secondary node, create a third action group.

```
virtual-link-loss-indication create action-group <third
action group>
```

36 At the secondary node, specify that the recovery action for the third action group is vc-mac-withdraw.

```
virtual-link-loss-indication add action-group <third
action group> type recovery action vc-mac-withdraw
precedence 1
```

Note: The vc-mac-withdraw message trigger by inverse VLLI (iVLLI) is supported only for static PWs.

37 At the secondary node, create an instance group (the first instance group)

```
virtual-link-loss-indication create instance-group
<first instance group>
```

38 At the secondary node, specify that the CFM service is the source for the first instance group.

```
virtual-link-loss-indication add instance-group <first
instance group> cfm-instance <cfm-service> mode source
```

39 At the secondary node, specify that the port-instance is the destination for the first instance group, and that the actions on the port-instance are those specified by the first action group. (The first action group specifies fault and recovery actions.)

```
virtual-link-loss-indication add instance-group <first
instance group> port-instance <port number> mode
destination action-group <first action group>
```

40 At the secondary node, enable the first instance group.

```
virtual-link-loss-indication enable instance-group
<first instance group>
```

41 At the secondary node, create a second instance group, specifying that it is bidirectional.

```
virtual-link-loss-indication create instance-group
<second instance group> direction bidirectional
```

Note: Bidirectionality allows failure and recovery actions to be propagated from source to destination, and from destination to source.

Copyright© 2022 Ciena® Corporation

42 At the secondary node, specify that the vc-instance is the source for the second instance group, and that the actions on the vc-instance are those specified by the third action group. (The third action group specifies a recovery action.)

```
virtual-link-loss-indication add instance-group <second
instance group> vc-instance <vc instance> mode source
action-group <third action group>
```

43 At the secondary node, specify that the port-instance is the destination for the second instance group, and that the actions on the port-instance are those specified by the second action group. (The second action group specifies fault and recovery actions.)

```
virtual-link-loss-indication add instance-group <second
instance group> port-instance <port number> mode
destination action-group <second action group>
```

44 At the secondary node, enable the second instance group.

```
virtual-link-loss-indication enable instance-group
<second instance group>
```

—end—

Copyright© 2022 Ciena® Corporation

Example

Commands at the primary node (node A):

```
port set port parent-port-at-primary y-protection primary


cfm enable
cfm service create vs vs-at-primary name cfm-service-at-primary level 4
cfm service enable service cfm-service-at-primary
cfm service set service cfm-service-at-primary alarm-priority 2 alarm-time 0
reset-time 100 y-protection-monitoring on
cfm mep create service cfm-service-at-primary sub-port sub-port-at-primary
type down mepid 1111


virtual-link-loss-indication create action-group A1
virtual-link-loss-indication add action-group A1 type fault action
port-unshut precedence 1
virtual-link-loss-indication add action-group A1 type recovery action
port-shut precedence 1
virtual-link-loss-indication create action-group a1
virtual-link-loss-indication add action-group a1 type fault action port-shut
precedence 1
virtual-link-loss-indication add action-group a1 type recovery action
port-unshut precedence 1
virtual-link-loss-indication create action-group a2
virtual-link-loss-indication add action-group a2 type recovery action
vc-mac-withdraw precedence 1
virtual-link-loss-indication create instance-group G1
virtual-link-loss-indication add instance-group G1 cfm-instance
cfm-service-at-primary mode source
virtual-link-loss-indication add instance-group G1 port-instance 1 mode
destination action-group A1
virtual-link-loss-indication enable instance-group G1
virtual-link-loss-indication create instance-group g1 direction bidirectional
virtual-link-loss-indication add instance-group g1 vc-instance <vc instance>
mode source action-group a2
virtual-link-loss-indication add instance-group g1 port-instance 1 mode
destination action-group a1
virtual-link-loss-indication enable instance-group g1
```


Copyright© 2022 Ciena® Corporation

Commands at the secondary node (node B):

```
port set port parent-port-at-secondary y-protection secondary


cfm enable
cfm service create vs vs-at-secondary name cfm-service-at-secondary level 4
cfm service enable service cfm-service-at-secondary
cfm service set service cfm-service-at-secondary alarm-priority 2 alarm-time
0 reset-time 100 y-protection-monitoring on
cfm mep create service cfm-service-at-primary sub-port sub-port-at-secondary
type down mepid 2222


virtual-link-loss-indication create action-group A1
virtual-link-loss-indication add action-group A1 type fault action
port-unshut precedence 1
virtual-link-loss-indication add action-group A1 type recovery action
port-shut precedence 1
virtual-link-loss-indication create action-group a1
virtual-link-loss-indication add action-group a1 type fault action port-shut
precedence 1
virtual-link-loss-indication add action-group a1 type recovery action
port-unshut precedence 1
virtual-link-loss-indication create action-group a2
virtual-link-loss-indication add action-group a2 type recovery action
vc-mac-withdraw precedence 1
virtual-link-loss-indication create instance-group G1
virtual-link-loss-indication add instance-group G1 cfm-instance
cfm-service-at-secondary mode source
virtual-link-loss-indication add instance-group G1 port-instance 1 mode
destination action-group A1
virtual-link-loss-indication enable instance-group G1
virtual-link-loss-indication create instance-group g1 direction bidirectional
virtual-link-loss-indication add instance-group g1 vc-instance <vc instance>
mode source action-group a2
virtual-link-loss-indication add instance-group g1 port-instance 1 mode
destination action-group a1
virtual-link-loss-indication enable instance-group g1
```


Copyright© 2022 Ciena® Corporation

Procedure 19-10
MPLS active/standby protection state management using VLLI in a Y-cable topology

Use this procedure to configure MPLS active/standby protection state management can using VLLI in a Y-cable topology.

The following figure shows the topology.

Figure 19-18
Y-cable topology for MPLS active/standby protection state management using VLLI

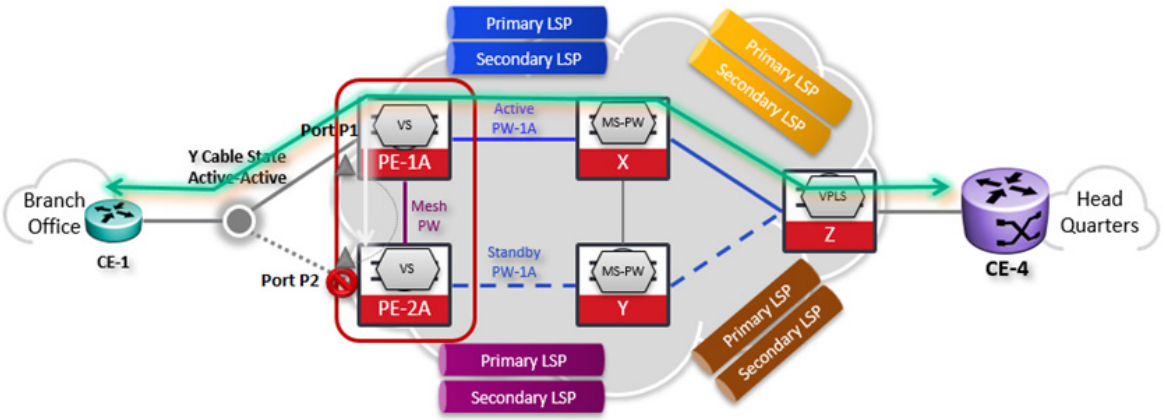

In the active/standby model, each split end of the Y-cable connects to a pseudo wire (PW), one being active and the other being standby. In the example, the Y-cable at the left connects to nodes PE-1A and PE-2A. At each of these nodes, the split end connects to a PW: active PW-1A at node PE-1A and standby PW-1A at node PE-2A.

The two nodes at the split ends are linked by a mesh PW. If a fault occurs, then traffic may or may not switch between the active PW-1A and the standby PW-1A. For example, if a fault occurs at the split-end port at node PE-1A then the split-end port at PE-2A would be used, and the mesh PW could carry traffic from node PE-2A to node PE-1A, which connects to the active PW-1A.

In this topology, the primary node is the node with the split-end port configured as y-protection primary. The secondary node is the node with the split-end port configured as y-protection secondary. VLLI decides which split-end port will be used. After startup or after a restart, VLLI selects the y-protection primary port. Subsequently VLLI changes ports if a fault requires it to do so, and auto-reversion does not occur.

Copyright© 2022 Ciena® Corporation

Prerequisites

At each node, primary and secondary, a sub-port has been created and has been attached to a virtual switch. You know the names of the sub-ports and the virtual switches, and you know the names of the parent ports, and their port numbers (to be used as the values of the “port-instance” parameter).

Step	Action
1	At the primary node, specify that the port connected to the split end of the Y-cable is the primary port. `port set port <port-name> y-protection primary`
2	At the primary node, enable CFM. `cfm enable`
3	At the primary node, create the CFM service on a virtual switch that uses the y-protected port. `cfm service create vs <vs-name> name <cfm-service> level <lvl>`
4	At the primary node, enable the CFM service. `cfm service enable service <cfm-service>`
5	For the CFM service, set the alarm priority, set the CFM alarm time, set the CFM service reset time, and enable y-protection monitoring. `cfm service set service <cfm-service> alarm-priority 2 alarm-time 0 reset-time 100 y-protection-monitoring on` ***Note 1:*** CFM service alarm time is the duration (in milliseconds) that CFM waits before sending a fault alarm to external clients (such as VLLI, G.8032, alarm engine). For Y-cable, the recommended alarm-time value is 0 milliseconds. The value should be configured in increments of 100 milliseconds. The default is three seconds. ***Note 2:*** CFM service reset time is the duration between the service fault clear event and the time that CFM notifies external clients (such as VLLI, G.8032, alarm engine). For Y-cable, the recommended reset time is 100 milliseconds. The default is three seconds. If the service fault occurs again during the reset time, there is no notification to external clients.
6	At the primary node, create the mep on the sub-port. `cfm mep create service <cfm-service> sub-port <sub-port-name> type up mepid <mepid>` where port-name is the name of the y-protected port

Copyright© 2022 Ciena® Corporation

Note: As shown in this procedure, you must set the y-protection monitoring parameter of a CFM service before creating the MEP. The y-protection monitoring parameter cannot be modified if a MEP exists for the service. If you want to modify the y-protection monitoring parameter of a CFM service, you must first delete all MEPs that have been created for that service.

7 At the primary node, create an action group (the first action group).

8 At the primary node, create an action group (the first action group).

```
virtual-link-loss-indication create action-group <first
action group>
```

9 At the primary node, specify that the fault action for the first action group is port-unshut.

```
virtual-link-loss-indication add action-group <first
action group> type fault action port-unshut precedence 1
```

10 At the primary node, specify that the recovery action for the first action group is port-shut.

```
virtual-link-loss-indication add action-group <first
action group> type recovery action port-shut precedence 1
```

11 At the primary node, create a second action group.

```
virtual-link-loss-indication create action-group <second
action group>
```

12 At the primary node, specify that the recovery action for the second action group is vc-mac-withdraw.

```
virtual-link-loss-indication add action-group <second
action group> type recovery action vc-mac-withdraw
precedence 1
```

Note: The vc-mac-withdraw message trigger by inverse VLLI (iVLLI) is supported only for static PWs.

13 At the primary node, create an instance group (the first instance group).

```
virtual-link-loss-indication create instance-group
<first instance group>
```

14 At the primary node, specify that the CFM service is the source for the first instance group.

```
virtual-link-loss-indication add instance-group <first
instance group> cfm-service <cfm service> mode source
```

15 At the primary node, specify that the port is the destination for the first instance group, and that the actions on the port are those specified by the first action group. (The first action group specifies fault and recovery actions.)

```
virtual-link-loss-indication add instance-group <first
instance group> port-instance <port number> mode
destination action-group A1
```

16 At the primary node, enable the first instance group.

```
virtual-link-loss-indication enable instance-group
<first instance group>
```


Copyright© 2022 Ciena® Corporation

17 At the primary node, create a second instance group.

```
virtual-link-loss-indication create instance-group
<second instance group>
```

18 At the primary node, specify that the port is the source for the second instance group.

```
virtual-link-loss-indication add instance-group <second
instance group> port-instance <port number> mode source
```

19 At the primary node, specify that the virtual circuit is the destination for the second instance group, and that the actions on the virtual circuit are those specified by the second action group. (The second action group specifies a recovery action.)

```
virtual-link-loss-indication add instance-group <second
instance group> vc-instance <vc instance> mode
destination action-group a2
```

20 At the primary node, enable the second instance group.

```
virtual-link-loss-indication enable instance-group
<second instance group>
```

21 At the secondary node, specify that the port connected to the split end of the Y-cable is the secondary port.

```
port set port <port-name> y-protection secondary
```

22 At the secondary node, enable CFM.

```
cfm enable
```

23 At the secondary node, create the CFM service on a virtual switch that uses the y-protected port.

```
cfm service create vs <vs-name> name <cfm-service> level
<lvl>
```

24 At the secondary node, enable the CFM service.

```
cfm service enable service <cfm-service>
```

25 For the CFM service, set the alarm priority, set the CFM alarm time, set the CFM service reset time, and enable y-protection monitoring.

```
cfm service set service <cfm-service> alarm-priority 2
alarm-time 0 reset-time 100 y-protection-monitoring on
```

Note 1: CFM service alarm time is the duration (in milliseconds) that CFM waits before sending a fault alarm to external clients (such as VLLI, G.8032, alarm engine). For Y-cable, the recommended alarm-time value is 0 milliseconds. The value should be configured in increments of 100 milliseconds. The default is three seconds.

Note 2: CFM service reset time is the duration between the service fault clear event and the time that CFM notifies external clients (such as VLLI, G.8032, alarm engine). For Y-cable, the recommended reset time is 100 milliseconds. The default is three seconds. If the service fault occurs again during the reset time, there is no notification to external clients.

Copyright© 2022 Ciena® Corporation

26 At the secondary node, create the mep on the sub-port.

```
cfm mep create service <cfm-service> sub-port
<sub-port-name> type up mepid <mepid>
```

where

port-name is the name of the y-protected port

Note: As shown in this procedure, you must set the y-protection monitoring parameter of a CFM service before creating the MEP. The y-protection monitoring parameter cannot be modified if a MEP exists for the service. If you want to modify the y-protection monitoring parameter of a CFM service, you must first delete all MEPs that have been created for that service.

27 At the secondary node, create an action group (the first action group).

```
virtual-link-loss-indication create action-group <first
action group>
```

28 At the secondary node, specify that the fault action for the first action group is port-unshut.

```
virtual-link-loss-indication add action-group <first
action group> type fault action port-unshut precedence 1
```

29 At the secondary node, specify that the recovery action for the first action group is port-shut.

```
virtual-link-loss-indication add action-group <first
action group> type recovery action port-shut precedence 1
```

30 At the secondary node, create a second action group.

```
virtual-link-loss-indication create action-group <second
action group>
```

31 At the secondary node, specify that the recovery action for the second action group is vc-mac-withdraw.

```
virtual-link-loss-indication add action-group <second
action group> type recovery action vc-mac-withdraw
precedence 1
```

Note: The vc-mac-withdraw message trigger by inverse VLLI (iVLLI) is supported only for static PWs.

32 At the secondary node, create an instance group (the first instance group).

```
virtual-link-loss-indication create instance-group
<first instance group>
```

33 At the secondary node, specify that the CFM service is the source for the first instance group.

```
virtual-link-loss-indication add instance-group <first
instance group> cfm-service <cfm service> mode source
```


Copyright© 2022 Ciena® Corporation

34 At the secondary node, specify that the port is the destination for the first instance group, and that the actions on the port are those specified by the first action group. (The first action group specifies fault and recovery actions.)

```
virtual-link-loss-indication add instance-group <first
instance group> port-instance <port number> mode
destination action-group G1
```

35 At the secondary node, enable the first instance group.

```
virtual-link-loss-indication enable instance-group
<first instance group>
```

36 At the secondary node, create a second instance group.

```
virtual-link-loss-indication create instance-group
<second instance group>
```

37 At the secondary node, specify that the port is the source for the second instance group.

```
virtual-link-loss-indication add instance-group <second
instance group> port-instance <port number> mode source
```

38 At the secondary node, specify that the virtual circuit is the destination for the second instance group, and that the actions on the virtual circuit are those specified by the second action group. (The second action group specifies a recovery action.)

```
virtual-link-loss-indication add instance-group <second
instance group> vc-instance <vc instance> mode
destination action-group a2
```

39 At the secondary node, enable the second instance group.

```
virtual-link-loss-indication enable instance-group
<second instance group>
```

—end—

Example

Commands at the primary node (node PE-1A):

```
port set port parent-port-at-primary y-protection primary


cfm enable
cfm service create vs vs-at-primary name cfm-service-at-primary level 4
cfm service enable service cfm-service-at-primary
cfm service set service cfm-service-at-primary alarm-priority 2 alarm-time 0
reset-time 100 y-protection-monitoring on
cfm mep create service cfm-service-at-primary sub-port sub-port-at-primary
type down mepid 1111


virtual-link-loss-indication create action-group A1
virtual-link-loss-indication add action-group A1 type fault action
port-unshut precedence 1
virtual-link-loss-indication add action-group A1 type recovery action
port-shut precedence 1
```


Copyright© 2022 Ciena® Corporation

```
virtual-link-loss-indication create action-group a2
virtual-link-loss-indication add action-group a2 type recovery action
vc-mac-withdraw precedence 1

virtual-link-loss-indication create instance-group G1
virtual-link-loss-indication add instance-group G1 cfm-service
cfm-service-at-primary mode source
virtual-link-loss-indication add instance-group G1 port-instance 1 mode
destination action-group A1
virtual-link-loss-indication enable instance-group G1

virtual-link-loss-indication create instance-group g1
virtual-link-loss-indication add instance-group g1 port-instance 1 mode
source
virtual-link-loss-indication add instance-group g1 vc-instance
<mesh-vc-instance> mode destination action-group a2
virtual-link-loss-indication enable instance-group g1
```

Commands at the secondary node (node PE-2A):

```
port set port parent-port-at-secondary y-protection secondary


cfm enable
cfm service create vs vs-at-secondary name cfm-service-at-secondary level 4
cfm service enable service cfm-service-at-secondary
cfm service set service cfm-service-at-secondary alarm-priority 2 alarm-time
0 reset-time 100 y-protection-monitoring on
cfm mep create service cfm-service-at-primary sub-port sub-port-at-secondary
type down mepid 2222


virtual-link-loss-indication create action-group A1
virtual-link-loss-indication add action-group A1 type fault action
port-unshut precedence 1
virtual-link-loss-indication add action-group A1 type recovery action
port-shut precedence 1

virtual-link-loss-indication create action-group a2
virtual-link-loss-indication add action-group a2 type recovery action
vc-mac-withdraw precedence 1

virtual-link-loss-indication create instance-group G1
virtual-link-loss-indication add instance-group G1 cfm-service <cfm service>
mode source
virtual-link-loss-indication add instance-group G1 port-instance 2 mode
destination action-group A1
virtual-link-loss-indication enable instance-group G1

virtual-link-loss-indication create instance-group g1
virtual-link-loss-indication add instance-group g1 port-instance 2 mode
source
virtual-link-loss-indication add instance-group g1 vc-instance
<mesh-vc-instance> mode destination action-group a2
virtual-link-loss-indication enable instance-group g1
```


Copyright© 2022 Ciena® Corporation

Procedure 19-11
Configuring VLLI to block a pseudowire

Configure VLLI to automatically block a pseudowire (PW) by triggering a pw-status-fault action when a port fault or CCM loss is detected. The following figure shows a sample scenario for shutting down a PWwhen VLLI detects CCM loss. This scenario is based on a CFM service running on a VLAN.

Figure 19-19
Fiber cut impacting connection between two switches

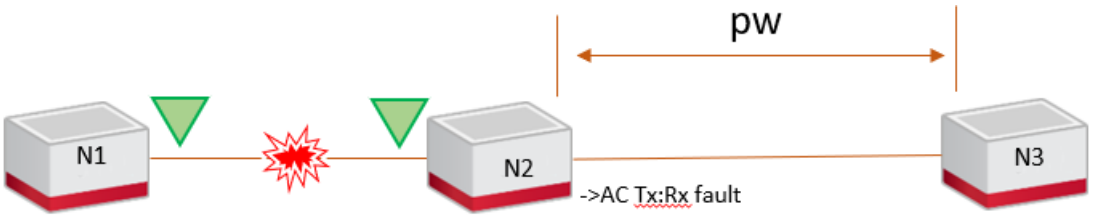

A failure in the link between N1 and N2 stops the CCM flow between two nodes hence raising a CFM fault at N2. This fault is detected by the VLLI and sends an action of pw-status-fault to the mpls L2vpnMgr which is then propagated using the link of N2-N3 blocking the PW from both the ends, that is N2 and N3.

As soon as the link between N1 and N2 heal, the CCMs are received by the node N2, healing the CFM service. The VLLI detects this recovery and sends pw-status-clear to the mpls L2vpnMgr, propagating the link N2-N3 unblocking the PW from both the ends, that is N2 and N3.

Step	Action
1	Enable CFM: `cfm enable`
2	Create a CFM service: `cfm service create vs <vs> name <string [45]> level <NUMBER: 0..7>`
3	Enable the CFM service created in step 2: `cfm service enable service <service>`
4	Create a MEP for the CFM service: `cfm mep create service <service> port <port>type up/down mepid <NUMBER: 1..8191>`

Copyright© 2022 Ciena® Corporation

5 Assign an alarm priority to the CFM service:

```
cfm service set service <service? alarm-priority <NUMBER:
1..5>
```

6 Create a VLLI instance group:

```
virtual-link-loss-indication create instance-group
<String[31]>
```

7 Create a VLLI action group:

```
virtual-link-loss-indication create action-group
<String[15}>
```

8 Modify the instance group created in step 5 to set the source mode action to the CFM service:

```
virtual-link-loss-indication add instance-group
<instance-group> mode source cfm-instance <service>
```

9 Add the pw actions in the Action group:

```
virtual-link-loss-indication add action-group
<actiongroup-name> t ype fault action pw-status-fault
precedence <1-8>
```

```
virtual-link-loss-indication add action-group
<actiongroup-name> t ype recovery action pw-status-clear
precedence <1-8>
```

10 Modify the instance group created in step 5 to set the destination mode to the PW:

```
virtual-link-loss-indication add instance-group
<instance-group> mode destination vc-instance <instance
name> action-group
```

11 Enable the instance group:

```
virtual-link-loss-indication enable instance-group
<instance-group>
```


Copyright© 2022 Ciena® Corporation

Procedure 19-12
Generating VLLI Fault and Recovery events based on pseudowires

This feature describes the support extended to pseudowires (PW) during fault and recovery scenarios in a VLLI configuration. In earlier scenarios, the support was only for ports impacting PW wherein multiple ports could impact a PW. In the current scenario, multiple source PW that monitors the fault can impact multiple target PW.

This feature supports the following mandatory and optional functions:

- Ability to cause a FAULT event if more than 1 instance within a designated group concurrently experiences a fault
- Fault on more than 1 Port Instances
- Fault on more than 1 PW Instances
- Ability for a FAULT event to trigger an action on a separate Instance
- Trigger an action on a PW Instance
- Trigger an action on a Port Instance
- Trigger an action on a CFM-service Instance (optional)
- Action of {pw-status-fault, pw-status-clear}
- Action of {port-shut, port-unshut}
- Action of {ccm-resume, ccm-stop} (optional)
- Cause a RECOVERY event if at least 1 Instance in a group is not faulted
- RECOVERY event to trigger an action on a separate Instance

Monitoring PWE in N:1 relationship

In the N:1 scenario, if there are multiple source PW being monitored that have a fault or recovery, then the corresponding and resulting target PW also gets impacted based on the fault or recovery scenario.

Copyright© 2022 Ciena® Corporation

Figure 19-20
Normal State

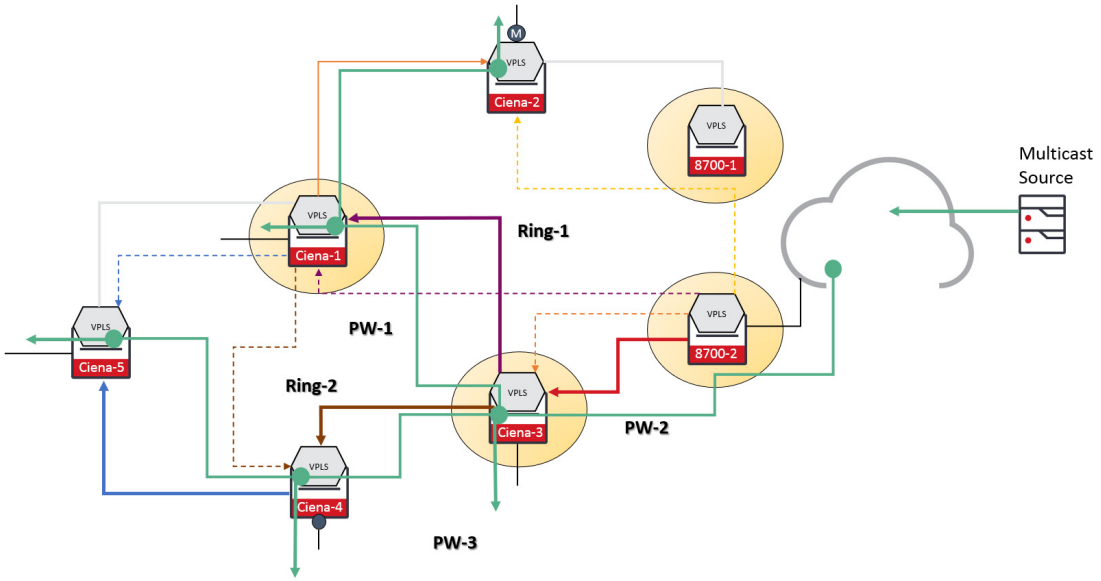

The above workflow describes a normal state where PW-1 and PW-2 are the monitored PWEs and operating without any fault or recovery. If there is a fault in PW-2, then the following workflow is initiated: *8700-2 > 8700-1 > Ciena-2 > Ciena-1 > Ciena-3 > Ciena-4, via PW-1 and PW-3.*

Copyright© 2022 Ciena® Corporation

Figure 19-21
Fault State

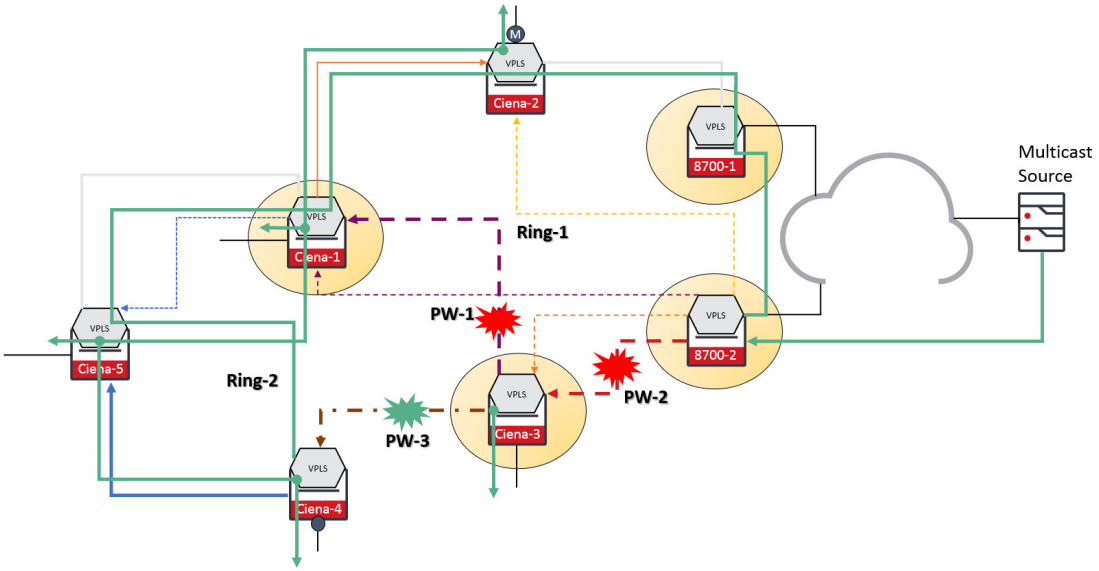

The above workflow describes the fault state where PW-1 and PW-2 are the monitored PWEs. If there is a fault detected in PW-1 and PW-2, then VLLI triggers a fault alert to the PW-3 target and subsequently to the Ciena-4. After Ciena-4 moves the PW-3 to Non-Forwarding and takes subsequent action for the next segment, then the following stream is initiated: *8700-2 > 8700-1> Ciena-2 > Ciena-1 > Ciena-5 > Ciena-4*.

Copyright© 2022 Ciena® Corporation

Figure 19-22
Recovery State

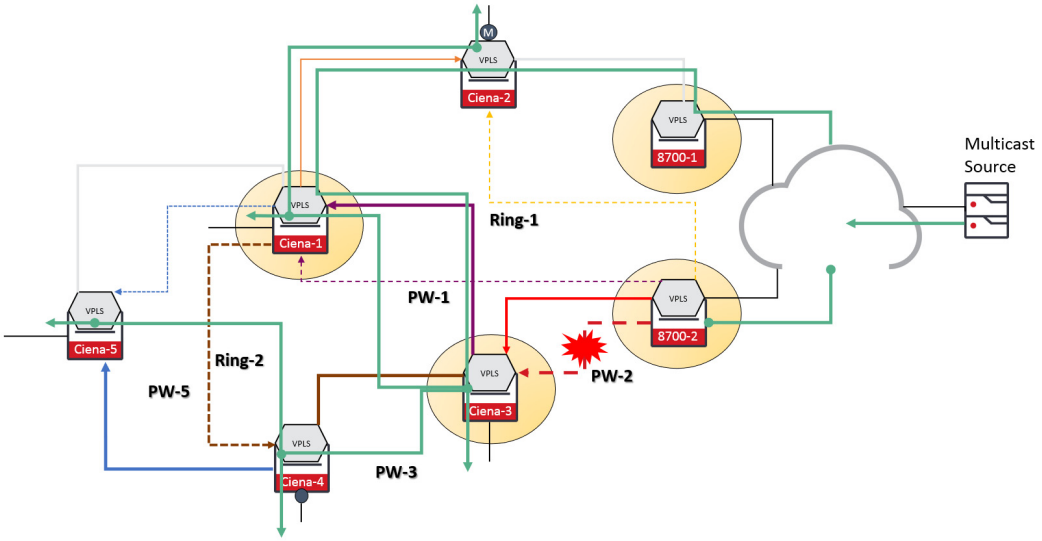

The above workflow describes the recovery state where PW-1 and PW-2 are the monitored PWEs and have recovered from a fault. After PW-1 or PW-2 recovers, the VLLI triggers a recovery alert to the PW-3 target and subsequently to the Ciena-4. After Ciena-4 moves the PW-3 to Forwarding state and takes subsequent action for the next segment, then one of the following streams is resumed based on the below recovery scenarios:

- **Scenario 1 (If PW-1 recovers)**: *8700-2 > 8700-1 > Ciena-2 > Ciena-1 > Ciena-3 > Ciena-4 via PW-1 and PW-3*.
- **Scenario 2 (If PW-2 recovers)**: *8700-2 > Ciena-3 > Ciena-4 via PW-2 and PW-3*.

Example

This requires monitoring of multiple PWE (PW-1 and PW-2) in a VPLS to propagate fault to single PWE (PW-3).

```
#PW Source Action Group // For Target PWE (PW-3)
virtual-link-loss-indication create action-group pw_src_act_grp
virtual-link-loss-indication add action-group pw_src_act_grp type fault
action pw-status-fault precedence 1
virtual-link-loss-indication add action-group pw_src_act_grp type recovery
action pw-status-clear precedence 1

#Instance Group // Bi-Directional Group // For PW-3
virtual-link-loss-indication create instance-group ig1 direction
bidirectional trigger all
virtual-link-loss-indication add instance-group ig1 vc-instance PW-1 mode
destination
virtual-link-loss-indication add instance-group ig1 vc-instance PW-2 mode
destination
```


Copyright© 2022 Ciena® Corporation

```
virtual-link-loss-indication add instance-group ig1 vc-instance PW-3 mode
source action-group pw_src_act_grp

virtual-link-loss-indication enable instance-group ig1
```

Note: While configuring the VLLI, you must set the monitored PWE (PW-1 and PW-2) as the destination despite being the source of fault.

Monitoring PWE in N:M relationship

In the N:M scenario, if there are multiple ports in the source PWEs that are being monitored and have a fault or recovery, then the corresponding and resulting target PW also gets impacted based on the fault or recovery scenario.

Note: The N:M relationship is not applicable for PWEs. Instead of PW-1 and PW-2, their underlying interfaces (Port-1 and Port-2) could be monitored.

Figure 19-23
Fault or Recovery State

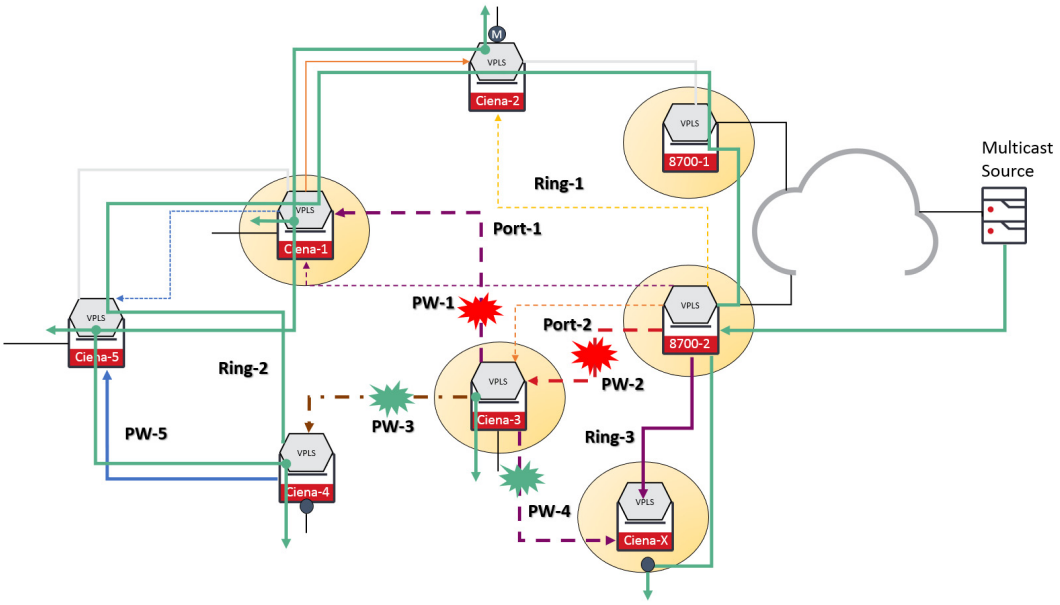

The above workflow describes the fault or recovery state where Port-1 and Port-2 are the underlying ports being monitored under PW-1 and PW-2.

- **Fault State**: If there is a fault detected on Port-1 and Port-2 of Ciena-3, then VLLI triggers a fault alert to the PW-3 and PW-4 targets. The following paths is applicable to this state:
  — *8700-2 > 8700-1 > Ciena-2 > Ciena-1 > Ciena-5 > Ciena-4*

Copyright© 2022 Ciena® Corporation

— *8700-2 > Ciena-X*

- **Recovery State**: After Port-1 and Port-2 of Ciena-3 recovers, PW-3 and PW-4 would also move to Forwarding state and takes subsequent action for the next segment. The following path is initiated in this state:

  — *8700-2 > 8700-1 > Ciena-2 > Ciena-1 > Ciena-3 > Ciena-4 (via PW-1 and PW-3)*

  — *8700-2 > 8700-1 > Ciena-2 > Ciena-1 > Ciena-3 > Ciena-X (via PW-1 and PW-4)*

Example

This requires monitoring of multiple Ports (Port-1 and Port-2) to propagate fault to multiple PWE (PW-3 and PW-4).

```
#PW Source Action Group // For Target PWE (PW-3, PW-4)
virtual-link-loss-indication create action-group pw_src_act_grp
virtual-link-loss-indication add action-group pw_src_act_grp type fault
action pw-status-fault precedence 1
virtual-link-loss-indication add action-group pw_src_act_grp type recovery
action pw-status-clear precedence 1

#Instance Group // Bi-Directional Group // For PW-3
virtual-link-loss-indication create instance-group ig1 direction
bidirectional trigger all
virtual-link-loss-indication add instance-group ig1 port-instance port 1 mode
destination
virtual-link-loss-indication add instance-group ig1 port-instance port 2 mode
destination
virtual-link-loss-indication add instance-group ig1 vc-instance PW-3 mode
source action-group pw_src_act_grp
virtual-link-loss-indication enable instance-group ig1

#Instance Group // Bi-Directional Group // For PW-4
virtual-link-loss-indication create instance-group ig2 direction
bidirectional trigger all
virtual-link-loss-indication add instance-group ig2 port-instance port 1 mode
destination
virtual-link-loss-indication add instance-group ig2 port-instance port 2 mode
destination
virtual-link-loss-indication add instance-group ig2 vc-instance PW-4 mode
source action-group pw_src_act_grp

virtual-link-loss-indication enable instance-group ig2
```

Note: A separate VLLI instance is required for each target PWE, even if the same set of ports are impacting them.

Output

Normal State

```
pnsw-14-7-5-5160*> virtual-link-loss-indication show instance-group ig1
+--------------- VIRTUAL-LINK-LOSS-INDICATION INSTANCE GROUP INFO ---------+
|      Name     | ig1                                                      |
```


Copyright© 2022 Ciena® Corporation

```
+--------------+-------------------------------------------------------------+
|   Parameter  |   Value                                                     |
+--------------+-------------------------------------------------------------+
| Admin Status | Enabled                                                     |
| Direction    | Bidirectional                                               |
| Trigger      |All                                                          |
+--------------+-------------------------------------------------------------+
+---------- VIRTUAL-LINK-LOSS-INDICATION INSTANCE INFO ----------+
| Type |Name            | Mode |  Action Group  |   Last Event   |
+------+----------------+------+----------------+----------------+
| VC   |PW-1            | Dest |                | Recovery       |
| VC   |PW-2            | Dest |                | Recovery       |
| VC   |PW-3            | Src  | pw_src_act_grp | Admin Recovery |
+------+----------------+------+----------------+----------------+

+------------------------------ PSEUDOWIRE TABLE ---------------+

|PW   | PW   | PW |       | Customer| Peer    | In     | Out    |      |

|Type| NH   | ID | Name | name    | IP      | Label | Label |Flags|

+----+-----+----+------+---------+--------+-------+-------+-----+

|ER  |1    |1   |PW-1  |         |4.4.4.4 |6004   |6004   |sUMoF|

|ER  |2    |2   |PW-2  |         |6.6.6.6 |6005   |6005   |sUMoF|

|ER  |3    |3   |PW-3  |         |3.3.3.3 |6006   |6006   |sUMoF|

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

Fault State

```
+--------------- VIRTUAL-LINK-LOSS-INDICATION INSTANCE GROUP INFO ---------+
|      Name    | ig1                                                         |
+--------------+-------------------------------------------------------------+
|   Parameter  |   Value                                                     |
+--------------+-------------------------------------------------------------+
| Admin Status | Enabled                                                     |
| Direction    | Bidirectional                                               |
| Trigger      |All                                                          |
+--------------+-------------------------------------------------------------+
+---------- VIRTUAL-LINK-LOSS-INDICATION INSTANCE INFO -------+
| Type |Name            | Mode |  Action Group  |   Last Event|
+------+----------------+------+----------------+-------------+
| VC   |PW-1            | Dest |                | Fault       |
| VC   |PW-2            | Dest |                | Fault       |
| VC   |PW-3            | Src  | pw_src_act_grp | Admin Fault |
+------+----------------+------+----------------+-------------+

+------------------------------ PSEUDOWIRE TABLE ---------------+

|PW   | PW   | PW |       | Customer| Peer    | In     | Out    |      |

|Type| NH   | ID | Name | name    | IP      | Label | Label |Flags|

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


Copyright© 2022 Ciena® Corporation

```
|ER  |1    |1   |PW-1  |         |4.4.4.4 |6004   |6004   |sDMoB|
|ER  |2    |2   |PW-2  |         |6.6.6.6 |6005   |6005   |sDMoB|
|ER  |3    |3   |PW-3  |         |3.3.3.3 |6006   |6006   |sDMoB|
+----+-----+-----------+---------+--------+-------+-------+-----+
```


Copyright© 2022 Ciena® Corporation

Service access control

Subscribers can connect a myriad of devices to the network. Keeping track of the changing list of these MAC addresses can be extremely difficult. SAOS devices employ Service Access Control (SAC) to

- allow static MAC addresses
- control the number of dynamically learned MAC addresses on a per-port basis

By configuring static MAC entries, security is provided to allow only those frames to be forwarded that have a source MAC address that is recognized by the device.

Once the maximum number of dynamically learned MAC addresses is reached, additional MAC addresses are not learned and frames from subsequent devices can be configured to either be forwarded or dropped. The learn limit does not apply to static MAC addresses.

For example, if the learn limit is set to four, the device learns the MAC addresses of the first four subscriber devices connected to the network and allow them to send frames, and can drop frames that come from a fifth device assuming that all devices are connected to the network at the same time. Frames containing source addresses that match static addresses in the device are forwarded, and are unaffected by the MAC limit set on the port. Although only four devices are connected to the network at one time, devices may be unplugged from the network and new devices, that is, new MAC addresses, are allowed once the old MAC addresses age out.

> ***Note:*** Limiting MAC addresses does not provide absolute security as a MAC address can still be spoofed, that is made to look like an authorized MAC address.

SAC protects the provider network from unauthorized access and can help prevent denial of service attacks. In addition, service providers can use SAC to offer different subscriber access packages without having to know any specifics about the subscriber devices. For example, a SAC can be set for

Copyright© 2022 Ciena® Corporation

"Gold" level subscribers that allows them to connect up to five devices (that is five MAC addresses) to the network, while "Bronze" level subscribers are only allowed to connect one device to the network at a time.

Procedures for SAC are:

- "Configuring SAC with MAC aging" on page 20-3
- "Configuring static MACs and SAC on the same device" on page 20-6
- "Configuring the device to forward or drop unknown frames" on page 20-10

Copyright© 2022 Ciena® Corporation

Procedure 20-1
Configuring SAC with MAC aging

Configure SAC with MAC aging. MAC aging is enabled by default.

To add a static MAC on a port, the VLAN must exist and the port must be a member of the VLAN.

Step	Action
1	Set the maximum number of dynamic MACs for the port. `flow access-control set port <PortName> {max-dynamic-macs <NUMBER: 0-16000>} [forward-unlearned <on|off>]` where port <PortName> — is the port name, expressed as a comma separated list or range. max-dynamic-macs <NUMBER: 0-16000> — is the MAC learning limit. The limit varies on a per-device basis. forward-unlearned <on|off> — • on: flood • off: drop
2	Enable MAC aging. `flow aging enable`
3	Set MAC aging. `flow aging set time <Seconds: 10-1000000>` where time <Seconds: 10-1000000> — is the number of seconds. The default value is 300 seconds.
4	Enter the static MAC address: `flow static-mac add vlan <VlanID> mac <xx:xx:xx:xx:xx:xx> port <PortNameList>}`

Copyright© 2022 Ciena® Corporation

5 (Optional) Verify the SAC configuration:

```
flow access-control show
```

6 (Optional) Verify the MAC aging settings:

```
flow aging show
```

7 (Optional) Verify the static MAC entry:

```
flow static-mac show
```

—end—

Example

This example shows sample output for the flow access-control show command.

```
flow access-control show
+------------------------- FLOW ACCESS-CONTROL TABLE --------------------------+
|    Instance     |     Dynamic Addresses      |  Forward  | Threshold |   Oper   |
| Type |  Value   | Limit  | Learnd | HWMark | Unlearned |  Reached  |  Status  |
+------+---------+--------+--------+--------+-----------+-----------+----------+
| Port |       1 |  32000 |      0 |      0 |     on    |     No    |  Enabled |
| Port |       2 |  32000 |      0 |      0 |     on    |     No    |  Enabled |
| Port |       3 |  32000 |      0 |      0 |     on    |     No    |  Enabled |
| Port |       4 |  32000 |      0 |      0 |     on    |     No    |  Enabled |
| Port |       5 |      4 |      0 |      0 |     on    |     No    |  Enabled |
| Port |       6 |  32000 |      0 |      0 |     on    |     No    |  Enabled |
| Port |       7 |  32000 |      0 |      0 |     on    |     No    |  Enabled |
| Port |       8 |  32000 |      0 |      0 |     on    |     No    |  Enabled |
| Port |       9 |  32000 |      0 |      0 |     on    |     No    |  Enabled |
| Port |      10 |  32000 |      0 |      0 |     on    |     No    |  Enabled |
| Port |      11 |  32000 |      0 |      0 |     on    |     No    |  Enabled |
| Port |      12 |  32000 |      0 |      0 |     on    |     No    |  Enabled |
| Port |    ag-1 |  32000 |      0 |      0 |     on    |     No    |  Enabled |
+------+---------+--------+--------+--------+-----------+-----------+----------+
```

This example shows sample output for the flow aging show command.

```
flow aging show
+------- FLOW AGING-INFO -------+
| Parameter        | Value       |
+------------------+-------------+
| Status           | Enabled     |
| Time (Seconds)   | 120         |
+------------------+-------------+
```


Copyright© 2022 Ciena® Corporation

This example shows sample output for the flow static-mac show command.

```
> flow static-mac show
+-------- FLOW STATIC-MAC TABLE --------+
| VLAN   | Address           | Port     |
+--------+-------------------+----------+
| 300    | AA:BB:CC:DD:EE:FF | 5        |
+--------+-------------------+----------+
```


Copyright© 2022 Ciena® Corporation

Procedure 20-2
Configuring static MACs and SAC on the same device

Configure static or dynamic MAC entries to restrict access to a service.

Make sure that one of these activities has been performed:

- NIC is configured with a static IP address that is in the subnet of the network
- NIC is configured to use DHCP to obtain an IP address and a DHCP server is reachable on the network

To add a static MAC on the port, the VLAN used must already be created and the port must be a member of the VLAN.

Step	Action
1	Set the maximum number of allowed MAC addresses. `flow access-control set port <port> {max-dynamic-macs <NUMBER>}` where port <port> — is the port. max-dynamic-macs <NUMBER> — is the maximum number of dynamic MACs allowed.
2	Enter the static MAC address. `flow static-mac add mac <MAC address: XX:XX:XX:XX:XX:XX> port <Port>} vlan <VLAN>` where {mac <MAC address: XX:XX:XX:XX:XX:XX>} — is the static MAC address. {port <Port>} — is the port. {vlan <VLAN>} — is the VLAN.
3	(Optional) Verify the SAC configuration. `flow access-control show`
4	(Optional) Verify the static MAC entry. `flow mac-addr show`

Copyright© 2022 Ciena® Corporation

5 Set the maximum number of allowed MAC addresses for the second port.

```
flow access-control set port <port> max-dynamic-macs
<NUMBER>
```

where

port <port>	is the port.
max-dynamic-macs <NUMBER>	is the maximum number of dynamic MACs allowed.

6 Disable MAC aging. MAC aging is disabled by default.

```
flow aging disable
```

7 (Optional) Verify the SAC configuration:

```
flow access control show
```

—end—

Example

In this example, a static MAC address is created on port 5 for the set top box, to limit access to video services for a specific set top box, along with disabling MAC learning on the port by setting an SAC entry with a learn limit of 0. In this example, only the set top box has access to the video services, and no other MACs are serviced on the port.

On a different port, that is, port 6, an SAC entry with a learn limit of 1 is set for the port so that only one device, with a dynamic address, can be learned and serviced by the port. In this example, access to a port is limited to one device, such as a set top box, for ports 5 and 6:

- On port 5, the device's MAC address is entered manually, no dynamic learning takes place on that port and no other devices are accepted on this port. On port 6, the MAC address is learned dynamically, but MAC address aging is disabled for the device.
- On port 6, if the Ethernet device is removed, its source MAC address remains in the MAC Learning table until the MAC table is flushed. Until such time, the SAOS device does not forward packets from a different device if it is attached to port 6.
- All other ports have the default SAC entry.

Set the maximum number of allowed MAC addresses to 0.

```
flow access-control set port 5 max-dynamic-macs 0
```


Copyright© 2022 Ciena® Corporation

Add a static MAC to port 5 on VLAN 300.

```
vlan create vlan 300
vlan add vlan 300 port 5
flow static-mac add vlan 300 mac aa:bb:cc:dd:ee:ff port 5
```

Verify the SAC configuration.

```
> flow access-control show

+------------------------- FLOW ACCESS-CONTROL TABLE --------------------------+
|     Instance      |      Dynamic Addresses      |  Forward  | Threshold |  Oper  |
| Type |  Value     | Limit  | Learnd | HWMark    | Unlearned | Reached   | Status |
+------+------------+--------+--------+-----------+-----------+-----------+--------+
| Port | 1          | 32000  |      0 |         0 |        on |        No |  Enabled |
| Port | 2          | 32000  |      0 |         0 |        on |        No |  Enabled |
| Port | 3          | 32000  |      0 |         0 |        on |        No |  Enabled |
| Port | 4          | 32000  |      0 |         0 |        on |        No |  Enabled |
| Port | 5          | 32000  |      0 |         0 |        on |        No |  Enabled |
| Port | 6          | 32000  |      0 |         0 |        on |        No |  Enabled |
| Port | 7          | 32000  |      1 |         1 |        on |        No |  Enabled |
| Port | 8          | 32000  |      0 |         0 |        on |        No |  Enabled |
| Port | 9          | 32000  |      0 |         0 |        on |        No |  Enabled |
| Port | 10         | 32000  |      0 |         0 |        on |        No |  Enabled |
| Port | 11         | 32000  |      0 |         0 |        on |        No |  Enabled |
| Port | 12         | 32000  |      0 |         0 |        on |        No |  Enabled |
+------+------------+--------+--------+-----------+-----------+-----------+--------+
```

Verify the static MAC entry.

```
> flow mac-addr show
+------------------------- FLOW MAC-LEARN TABLE -------------------------+
| VLAN    | Address (SA)       | Port               | Type  | Virtual Switch |
+---------+--------------------+--------------------+-------+----------------+
| 1       | 00:02:A1:30:A8:1D  | 7                  | Dyna  | ---            |
| 1       | 00:03:18:6D:A9:62  | 1                  | Stat  | ---            |
| 1       | 00:03:18:6D:A9:63  | 2                  | Stat  | ---            |
| 1       | 00:03:18:6D:A9:64  | 3                  | Stat  | ---            |
| 1       | 00:03:18:6D:A9:65  | 4                  | Stat  | ---            |
| 1       | 00:03:18:6D:A9:66  | 5                  | Stat  | ---            |
| 1       | 00:03:18:6D:A9:67  | 6                  | Stat  | ---            |
| 1       | 00:03:18:6D:A9:68  | 7                  | Stat  | ---            |
| 1       | 00:03:18:6D:A9:69  | 8                  | Stat  | ---            |
| 1       | 00:03:18:6D:A9:6A  | 9                  | Stat  | ---            |
| 1       | 00:03:18:6D:A9:6B  | 10                 | Stat  | ---            |
| 1       | 00:03:18:6D:A9:6C  | 11                 | Stat  | ---            |
| 1       | 00:03:18:6D:A9:6D  | 12                 | Stat  | ---            |
+---------+--------------------+--------------------+-------+----------------+
```

Set the maximum number of allowed MAC addresses for the second port to 1.

```
flow access-control set port 6 max-dynamic-macs 1
```

Disable MAC aging.

```
flow aging disable
```


Copyright© 2022 Ciena® Corporation

Verify the SAC configuration.

```
> flow access-control show
+------------------------- FLOW ACCESS-CONTROL TABLE --------------------------+
|     Instance      |     Dynamic Addresses      |  Forward  | Threshold |  Oper  |
| Type |  Value     | Limit  | Learnd | HWMark | Unlearned | Reached   | Status |
+------+------------+--------+--------+--------+-----------+-----------+--------+
| Port | 1          | 32000  |      0 |      0 |      on   |      No   | Enabled |
| Port | 2          | 32000  |      0 |      0 |      on   |      No   | Enabled |
| Port | 3          | 32000  |      0 |      0 |      on   |      No   | Enabled |
| Port | 4          | 32000  |      0 |      0 |      on   |      No   | Enabled |
| Port | 5          | 32000  |      0 |      0 |      on   |      No   | Enabled |
| Port | 6          | 32000  |      0 |      0 |      on   |      No   | Enabled |
| Port | 7          | 32000  |      1 |      1 |      on   |      No   | Enabled |
| Port | 8          | 32000  |      0 |      0 |      on   |      No   | Enabled |
| Port | 9          | 32000  |      0 |      0 |      on   |      No   | Enabled |
| Port | 10         | 32000  |      0 |      0 |      on   |      No   | Enabled |
| Port | 11         | 32000  |      0 |      0 |      on   |      No   | Enabled |
| Port | 12         | 32000  |      0 |      0 |      on   |      No   | Enabled |
+------+------------+--------+--------+--------+-----------+-----------+--------+
```


Copyright© 2022 Ciena® Corporation

Procedure 20-3
Configuring the device to forward or drop unknown frames

The Forward-Unlearned option allows for the forwarding or dropping of frames containing unknown, that is, unlearned, source MAC addresses when the SAC limit has been reached on a port.

Once the maximum number of dynamic MACs has been learned on a port, subsequent frames received with an unlearned source MAC are not learned and the frames are forwarded or dropped based on the setting of the Forward-Unlearned option. The Forward-Unlearned option is turned on by default for each port.

Step	Action

To turn the Forward-Unlearned option on

1 Turn the Forward-Unlearned option on:

```
flow access-control set port <port> forward-unlearned on
```

where

port <port>	is the port.
forward-unlearned on	sets the Forward-Unlearned option to on.

To turn the Forward-Unlearned option off

2 Turn the Forward-Unlearned option off:

```
flow access-control set port <port> forward-unlearned off
```

where

port <port>	is the port.
forward-unlearned off	sets the Forward-Unlearned option to off.

—end—

Copyright© 2022 Ciena® Corporation

Flow loop detection

MAC learning lets the system update source port information if MAC SA of the incoming frame is known. The frame's ingress port also does not match the port information in the L2 entry in the MAC database. This requires the address to be relearned due to station movement. This capability can be leveraged to detect continuous MAC flapping between two or more ports which indicates a network loop exists.

Flow loop detection catches loops existing in external (customer or third party) networks. Loops in a network are detected by 802.1ag Connectivity Fault Management (CFM). Control protocol (such as xSTP and CFM) based loop detection is not practical on UNI or E-NNI interfaces because there is more than one administrator involved in control protocol provisioning.

Customer network loops cause MAC flapping in provider networks as shown in this figure.

Copyright© 2022 Ciena® Corporation

Figure 21-1
MAC flapping caused by UNI-C looping

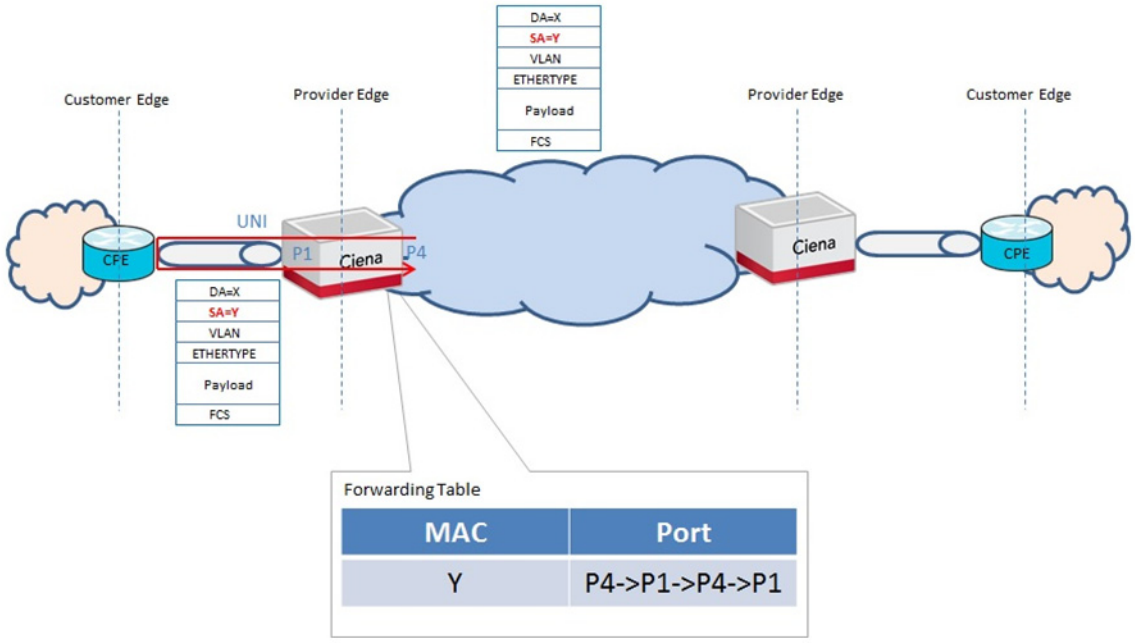

This MAC flapping does not allow the provider network to settle and inject BUM frames in the provider network. Broadcast containment may not be enough for aggressive ingress loops. Networks can be protected by detecting network loops and preventing them by configuring loop-terminating action at edge devices.

Copyright© 2022 Ciena® Corporation

Loop detection

Flow loop detection is built over the system's dynamic MAC learning capability. Loops are detected based on user-configurable MAC motion count. The system keeps track of MAC movements and if any MAC goes through MAC motion count, the number of movements in the system's detection interval, the system declares a loop between two or more ports.

MAC movements are counted in software based on hardware notifications, and may not match the true MAC movements in the hardware.

MAC movement based loop detection catches loops over these flows:

- Provider network > customer network > provider network
- Customer network > provider network > customer network

This figure shows the provider network > customer network > provider network flow.

Figure 21-2
Provider network flow

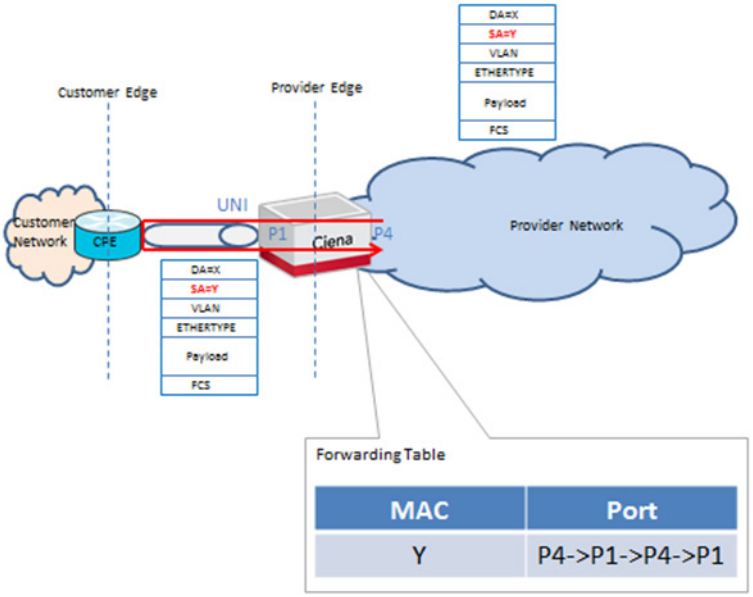

Copyright© 2022 Ciena® Corporation

This figure shows the customer network -> provider network -> customer network flow.

Figure 21-3
Customer network flow

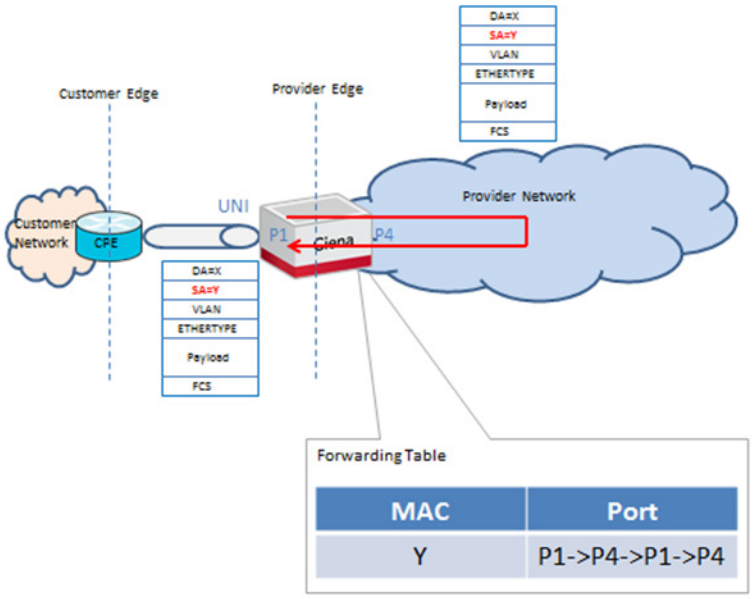

Flow loop detection is ideally used to detect loops caused by the provider network -> customer network -> provider network only. For the customer network -> provider network -> customer network flow, it is assumed that the provider network is running a network loop prevention mechanism such as xSTP, G.8032, or VPLS to avoid loops in the network.

> ***Note 1:*** Flow loop detection is supported on any port, but it is recommended to use it for UNI or E-NNI ports. For example, in either of the previously mentioned flows, it is recommended to enable flow loop detection on port P1.
>
> ***Note 2:*** Flow loop detection does not benefit UNI ports configured as an attachment circuit of an MPLS service.

Loop notification

When the system discovers an ingress loop on a UNI or E-NNI, it notifies the user by generating SNMP traps and logging events. In the case of notify0only, only notifications are sent to the user. Notifications are sent even if the configured actions are Deny MAC movement or Port shutdown.

Copyright© 2022 Ciena® Corporation

Loop prevention

When the system discovers an ingress loop on a UNI or E-NNI, the system tries to prevent a loop by automatically initiating these loop prevention actions on that port:

- Deny MAC movement — Disallow MAC movement on the port and drop all frames attempting to trigger a MAC movement.
- Port shutdown — Operationally bring down the port forcibly

The system supports auto-reversion of these loop-prevention actions from the UNI or E-NNI at the expiry of a user-configurable hold-off timer if the loop condition has been resolved, meaning the UNI or E-ENNI is no longer receiving frame(s) attempting to trigger MAC movement.

The system determines the loop state on the UNI or E-NNI without reverting to “Deny MAC movement” action. In the case of a port shutdown action, the system cannot determine whether looping occurs after bringing up the port. It proceeds with the assumption that the loop was cleared.

Auto-reversion can be turned off at the port level if manual intervention is desired in removing the loop preventing actions.

> ***Note 1:*** The flow loop detection is desirable on UNI or E-NNI interfaces. The system does not prevent the administrator from using this capability on NNI ports.
>
> ***Note 2:*** Frames dropped by “Deny MAC movement” action are counted by “Rx Discard Packets statistics at the ingress port.

Procedures for flow loop detection are:

- “Enabling and disabling flow loop detection” on page 21-6
- “Enabling MAC learning at a virtual switch or VLAN” on page 21-7
- “Configuring flow loop detection” on page 21-8
- “Resetting flow loop detection to default values” on page 21-10
- “Clearing loop detection statistics on the specified port(s)” on page 21-11
- “Displaying flow learning information” on page 21-12
- “Displaying flow loop detection information” on page 21-13

Copyright© 2022 Ciena® Corporation

Procedure 21-1
Enabling and disabling flow loop detection

You can enable flow loop detection

- globally
- on a UNI port(s)

You can also disable flow loop detection for the device or a specified port(s).

> ***Note:*** If the virtual switch includes more than one port member, flow loop detection must be on enabled on all UNIs.

Step	Action
Globally enable flow loop detection using MAC movement	
1	Globally enable loop detection: `flow loop-detection enable`
Enable flow loop detection on UNI ports	
2	Enable flow loop detection on UNI ports: `flow loop-detection enable port [<port>]` where port <port> is the UNI port you are enabling flow loop detection functionality.
Disable flow loop detection for the device or a specified port(s)	
3	Disable flow loop detection for the device or a specified port(s): `flow loop-detection disable [<port>]`

—end—

Example

This example enables flow loop detection on port 1.

```
flow loop-detection enable port 1
```

This example pertains to a virtual switch that includes more than one port member.

```
flow loop-detection enable port 1
flow loop-detection enable port 2
flow loop-detection enable port 3
```


Copyright© 2022 Ciena® Corporation

Procedure 21-2
Enabling MAC learning at a virtual switch or VLAN

Make sure that MAC learning at a virtual switch or VLAN is enabled. By default, MAC learning is enabled at the virtual switch and the VLAN. If it is not enabled, enable it.

Step	Action
1	Make sure that MAC learning is enabled on a virtual switch or VLAN: `flow learning enable vlan <vlan name>` where vlan <vlan name> is the VLAN MAC learning is enabled on.

—end—

Example

This example enables MAC learning on vlan vlan1.

```
flow learning enable vlan vlan1
```


Copyright© 2022 Ciena® Corporation

Procedure 21-3
Configuring flow loop detection

You can configure flow loop detection:

- Optionally configure loop prevention action on UNI port(s)
- Optionally configure the MAC-motion count
- Disable the reversion of the loop detection action if manual intervention is desired to remove the loop prevention action
- Optionally configure the hold-off timer

Step	Action

Configure loop prevention action on UNI port(s)

1 Optionally configure loop prevention action on UNI port(s):

```
flow loop-detection set port <port> action <block-mac-
movement | notify-only|port-shut>
```

where

port <port>	is the UNI port you are setting the loop prevention action on.
action <block-mac-movement | notify-only | port-shut>	sets the loop prevention action on the specified port(s). The default is block-mac-movement. block-mac movement drops frames attempting to perform MAC movement from another port to this port. port-shut is the admin down port. notify-only sends notification (SNMP traps/syslog event) only.

Configure MAC-motion count

2 Optionally configure MAC-motion count:

```
flow loop-detection set mac-motion-count <2..10>
```

where

mac-motion-count <2..10>	denotes the number of times a MAC address can flap across two or more ports to declare a loop condition. If any MAC observes MAC-motion count movements in the system's detection-interval across two or more ports, the system detects a loop between the participating ports.

Copyright© 2022 Ciena® Corporation

Configure auto reversion of loop prevention action on the specified port(s)

3 Set the auto reversion of loop prevention action on the specified port(s):

```
flow loop-detection set port <port> revertive <on|off>
```

where

port <port>	configures the flow loop detection on a port-by-port basis.
revertive <on|off>	turns the reversion of loop detection action on or off. In the case of off, the applied action on the port is not reverted. Default is on.

4 Revoke the running action from a port when a loop is found and the configured action is applied on a port. Once the feature is disabled from a port, flow loop detection action is automatically revoked.

```
flow loop-detection disable port <port>
```

Configure the hold-off time

5 Optionally configure the hold-off time for reversion to start:

```
flow loop-detection set reversion-holdoff-time <1..12>
```

where

reversion-holdoff-time <1..12>	is the loop prevention action hold-off time. In the case of auto reversion on, the applied loop prevention action on the port is reverted once the hold-off timer expires. Default is 5 minutes.

—end—

Example

This example configures flow loop prevention action on port 1 and brings down the port forcibly once the loop is detected on this port.

```
flow loop-detection set port 1 action port-shut
```


Copyright© 2022 Ciena® Corporation

Procedure 21-4
Resetting flow loop detection to default values

You can reset flow loop detection using attributes to default values:

- global MAC motion threshold count to the system default. Default is 4.
- global auto reversion hold-off time to the system default. Default is 5 minutes.

Step	Action

Reset the global MAC motion threshold count to the system default

1 Reset the global MAC motion threshold count to the system default:

```
flow loop-detection unset mac-motion count
```

Reset the global auto reversion hold-off time to the system default

2 Reset the global auto reversion hold-off time to the system default:

```
flow loop-detection unset reversion-holdoff-time
```

—end—

Copyright© 2022 Ciena® Corporation

Procedure 21-5
Clearing loop detection statistics on the specified port(s)

Clear loop detection statistics on the specified port(s).

Step	Action
1	Clear loop detection statistics on a specified port(s): `flow loop-detection clear port <port> statistics` **—end—**

Copyright© 2022 Ciena® Corporation

Procedure 21-6
Displaying flow learning information

You can display flow learning information.

Step	Action
1	Display flow learning information: `flow learning show all`

—end—

Example

This example shows output for the flow learning show all command.

```
> flow learning show all

+------------------------- VLAN MAC LEARNING -------------------------+
|                   VLAN                      |      Learning Status      |
|  ID  | Name                                 |    Admin     |    Oper    |
+------+--------------------------------------+--------------+------------+
|    1 | Default                              |   Enabled    |  Enabled   |
|   14 | VLAN#14                              |   Disabled   |  Disabled  |
|   15 | VLAN#15                              |   Enabled    |  Enabled   |
|  101 | VLAN#101                             |   Disabled   | VS Override |
|  102 | VLAN#102                             |   Enabled    | VS Override |
|  127 | Mgmt                                 |   Enabled    |  Enabled   |
|  201 | VLAN#201                             |   Enabled    |  Enabled   |
| 4001 | VLAN#4001                            |   Enabled    |  Enabled   |
| 4002 | VLAN#4002                            |   Enabled    |  Enabled   |
+------+--------------------------------------+--------------+------------+

+---------------- VS MAC LEARNING ----------------+
|                  | Active |   Learning Status   |
| VS               |  VLAN  |  Admin   |   Oper   |
+------------------+--------+----------+----------+
| vsE1             |   101  | Disabled | Disabled |
| vsE2             |   102  | Enabled  | Enabled  |
+------------------+--------+----------+----------+
```


Copyright© 2022 Ciena® Corporation

Procedure 21-7
Displaying flow loop detection information

You can display

- flow loop detection information on a specific port
- all flow loop detection information

Step	Action

Display flow loop detection information on a specific port

1 Display flow loop detection information on a specific port:

```
flow loop-detection show [<port>]
```

Display all flow loop detection information

2 Display all flow loop detection information:

```
flow loop-detection show
```

—end—

Example

This example shows the output for the flow loop-detection on port 1:

```
> flow loop-detection show port 1
+---------------- PORT 1 INFO ----------------+
| Parameter                 | Value            |
+---------------------------+------------------+
| Admin State               | Disabled         |
| Operational State         | Disabled         |
| Auto Reversion            | On               |
| Loop Prevent Action       | block-mac-motion |
| Loop Present              | No               |
| Loop Occurrence Count     | 0                |
+---------------------------+------------------+
```


Copyright© 2022 Ciena® Corporation

This example shows the output for the flow loop-detection show command:

```
> flow loop-detection show
+------------------- LOOP-DETECTION GLOBAL CONFIGURATION -------------------+
| Parameter                                | Value                            |
+------------------------------------------+----------------------------------+
| Admin State                              | Disabled                         |
| Motion Count                             | 4                                |
| Reversion Hold-Off Time(Minutes)         | 5                                |
+------------------------------------------+----------------------------------+

+-------------------------------------------------------------------------------+
| Port Table |           Admin Config                   |  Oper Status        |
+------------+----------+-----------+------------------+----------+----------+
| Port       | Admin    | Auto      | Loop Prevent     | Oper     | Loop     |
| Name       | State    | Reversion | Action           | Status   | Present  |
+------------+----------+-----------+------------------+----------+----------+
| 1          | Disabled |     On    | block-mac-motion | Disabled | No       |
| 2          | Disabled |     On    | block-mac-motion | Disabled | No       |
| 3          | Disabled |     On    | block-mac-motion | Disabled | No       |
| 4          | Disabled |     On    | block-mac-motion | Disabled | No       |
+------------+----------+-----------+------------------+----------+----------+
```


Copyright© 2022 Ciena® Corporation

Error codes

This section lists error codes.

Traffic profiling error codes

The Traffic profiling error codes table 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 overlaps with a traffic profile on the same port that has the same attribute value for A.

This example shows 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 22-1
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.

Copyright© 2022 Ciena® Corporation

Table 22-1
Traffic profiling error codes

4	Vlan^{A}	Profile with Vlan^{A} classifier. Profile with Vlan^{A} + any COS (IP, dscp, dot1d) classifiers.
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.

Copyright© 2022 Ciena® Corporation

Table 22-1
Traffic profiling error codes

19	Vlan[A] + IP[B]	Profile with IP[B] classifier and IP value matches.
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© 2022 Ciena® Corporation

Table 22-1
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© 2022 Ciena® Corporation

Table 22-1
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.

Copyright© 2022 Ciena® Corporation

Table 22-1
Traffic profiling error codes

50	VS^A + $dscp^B$	Profile with Vlan +$dscp^B$ classifiers and Vlan is a member of VS^A.
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.

Copyright© 2022 Ciena® Corporation

Table 22-1
Traffic profiling error codes

69	VS^A + $Vlan^B$ + IP^C	Profile with VS^A + $Vlan^B$ classifiers. Profile with VS^A + $Vlan^B$ + dot1d or dscp classifiers.
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© 2022 Ciena® Corporation

Copyright© 2022 Ciena® Corporation

39XX/51XX Switches and Platforms

Advanced Ethernet Configuration

Copyright© 2022 Ciena® Corporation. All rights reserved.

SAOS 6.21.2
Publication: 009-3458-040
Document status: Standard
Revision A
Document release date: March 2022

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