



39XX/51XX Service Delivery, Aggregation and Virtualization Switches

MPLS Configuration

SAOS 6.18.1

What's inside...

New in this release

MPLS configuration fundamentals

MPLS overview task flow

Initial configuration

Routing protocol configuration

RSVP-TE configuration

Label range configuration

Static tunnel configuration

Dynamic tunnel configuration

Additional tunnel configuration

LDP configuration

Virtual circuit configuration

Virtual switch configuration

L2 VPN service configuration

Interface configuration

MPLS troubleshooting

Bidirectional Forwarding Detection configuration

Alarm Indication Signal with Link Down Indication

Virtual private LAN service integrated routing and bridging

Seamless MPLS configuration

Flow loop detection for MPLS services

Loop-free alternate fast reroute configuration

Sample MPLS topology

009-3313-041 - Standard Revision A

June 2019

Copyright© 2019 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© 2019 Ciena® Corporation. All Rights Reserved

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

Security

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

Contacting Ciena

Corporate Headquarters	410-694-5700 or 800-921-1144	www.ciena.com
Customer Technical Support/Warranty		
In North America	1-800-CIENA-24 (243-6224) 410-865-4961	
In Europe, Middle East, and Africa	800-CIENA-24-7 (800-2436-2247) +44-207-012-5508 00 0800 77 454 (Slovenia)	
In Asia-Pacific	800-CIENA-24-7 (800-2436-2247) +81-3-6367-3989 +91-124-4340-600 120 11104 (Vietnam) 000 8004401369 (India)	
In Caribbean and Latin America	800-CIENA-24-7 (800-2436-2247) 1230-020-0845 (Chile) 009 800-2436-2247 (Colombia) 0800-77-454 (Mexico and Peru) 00 008000442510 (Panama)	
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 or (2) as a standalone application, and Software includes upgrades acquired by you from Ciena or a Ciena authorized reseller. 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 and only for its intended use as evidenced by the applicable product documentation. Unless the context does not permit, Software also includes associated documentation.

2. Open Source and Third Party Licenses. Software excludes any open source or third-party programs supplied by Ciena under a separate license, and you agree to be bound by the terms of any such license. If a separate license is not provided, any open source and third party programs are considered "Software" and their use governed by the terms of this License.

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.

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 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 period 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.

Publication history

June 2019

Standard Revision A.

First standard release of this document for SAOS 6.18.1.

Contents

About this document	xix
New in this release	1-1
Changes in Revision A	1-1
MPLS configuration fundamentals	2-1
MPLS layer 2 service architecture	2-3
VPWS	2-5
VPLS	2-6
Hierarchical VPLS	2-7
VPLS membership and MAC learning	2-8
MAC withdraw	2-9
MAC withdraw triggered by inverse VLLI in a Y-cable topology	2-10
MAC withdraw triggered by inverse VLLI in the active/active model	2-10
MAC withdraw triggered by inverse VLLI in the active/standby model	2-11
Routing protocols	2-12
Static routing	2-12
Dynamic routing	2-13
MPLS label operations	2-21
MPLS labels	2-23
Label Distribution Protocol	2-25
MPLS TP-TE gateway	2-25
Encapsulation	2-28
MEF8 encapsulation	2-28
MPLS encapsulation	2-28
Transport tunnels	2-30
Static unidirectional	2-30
Static associated bidirectional	2-31
Static co-routed bidirectional	2-32
Dynamic unidirectional	2-33
Dynamic associated bidirectional	2-33
Dynamic co-routed bidirectional	2-34
Traffic engineering	2-35
Resource Reservation Protocol - Traffic Engineering	2-35
Virtual circuits	2-36
Raw and tagged pseudowire type for virtual circuits	2-36
Pseudowire status signaling	2-38
Virtual circuit connectivity verification profile configuration	2-39

Pseudowire reversion	2-40
Multi-segment pseudowire	2-42
Forwarding Equivalence Class 129	2-45
Control word for pseudowire	2-46
Headers	2-46
Control word support	2-48
VCCV support	2-51
Limitations	2-52
Troubleshooting	2-52
LSP ping	2-52
LSP traceroute	2-53
VCCV ping	2-53
VCCV traceroute	2-53
Remote management for MPLS	2-54
Bidirectional Forwarding Detection	2-55
BFD sessions	2-57
Maximum number of BFD sessions	2-58
BFD mode	2-59
BFD roles	2-60
BFD control packet format	2-60
BFD state machine	2-61
Single-hop IPv4 BFD	2-62
Associated bidirectional LSPs	2-63
Interoperability	2-65
VCCV BFD	2-65
BFD dampening	2-66
Alarm Indication Signal with Link Down Indication	2-73
Operation	2-73
Remote Defect Indication	2-74
Multiple fault conditions	2-76
Node protection	2-77
AIS/LDI and client/server layers	2-78
Interworking with other vendor products	2-79
Features not available on 393x and 3916 platforms	2-80

Initial configuration **3-1**

List of procedures

3-1	Installing the MPLS license on 39XX/51XX	3-2
3-2	Configuring IP for MPLS interfaces on 39XX/51XX	3-4
3-3	Disabling RSTP and MSTP	3-6

MPLS overview task flow **4-1**

Additional tunnel configuration	4-5

Routing protocol configuration **5-1**

- SNMP support 5-2
 - OSPF routing protocol 5-2
 - IS-IS routing protocol 5-2

List of procedures

- 5-1 Configuring OSPF routing protocol 5-4
 - 5-2 Configuring IS-IS routing protocol 5-8
 - 5-3 Configuring static routes with numbered or unnumbered IP interfaces 5-18
 - 5-4 Configuring IS-IS on a point-to-point or unnumbered interface 5-19
 - 5-5 Configuring IS-IS route preference 5-20
-

RSVP-TE configuration **6-1**

- RSVP-TE rapid retry 6-2

List of procedures

- 6-1 Configuring RSVP-TE 6-4
-

Label range configuration **7-1**

List of procedures

- 7-1 Configuring label ranges 7-2
 - 7-2 Displaying label ranges 7-7
-

Static tunnel configuration **8-1**

- Next-hop diversity 8-1
- Static MPLS-TP co-routed tunnel support over the unnumbered IP interface 8-2
 - Tunnel FEC for static LSP 8-9

List of procedures

- 8-1 Configuring static TE tunnels 8-10
 - 8-2 Configuring static transit unidirectional TP tunnels 8-12
 - 8-3 Configuring static unidirectional ingress TP tunnels 8-13
 - 8-4 Configuring static unidirectional egress TP tunnels 8-15
 - 8-5 Configuring static bidirectional ingress-associated TE tunnels 8-16
 - 8-6 Configuring co-routed TP tunnels 8-17
 - 8-7 Displaying MPLS TE-tunnel information 8-24
 - 8-8 Displaying TP tunnel information 8-28
 - 8-9 Modifying tunnel attributes 8-33
 - 8-10 Inserting nodes in an unnumbered network 8-35
 - 8-11 Moving MPLS tunnels from numbered to unnumbered IP interfaces 8-39
-

Dynamic tunnel configuration **9-1**

- Resiliency and robustness 9-1
 - BFD linkage to IS-IS 9-2
 - Auto-backup 9-2
 - Make-before-break 9-3
 - LSP re-optimization 9-6
 - Graceful restart helper mode 9-7
 - Hardware-based IP forwarding 9-10
-

TE link metric	9-10
Resource affinity	9-11
SRLGs	9-12
Bandwidth management and reservation	9-13
Auto-bandwidth	9-13
LSP bandwidth auto-resize modes	9-14
MPLS QoS	9-15
DiffServ-TE	9-16
Class type	9-16
Bandwidth constraint	9-17
Queues and queue groups	9-18
Labels	9-19
MPLS fast reroute	9-19
Facility backup method	9-22
Link and node protection	9-22
Bandwidth protection	9-23
Link affinity	9-23
Setup and hold priority	9-24
Fast reroute protection switching	9-24
Fast reroute reversion	9-24
Fast reroute reselection	9-24
OAM support	9-26
Preemption	9-26
Soft preemption	9-27
Bandwidth allocation in DiffServ	9-27
Modes of operation	9-28
Preemption policy	9-32
Preemption for facility-bypass tunnels	9-32
Configuration examples	9-32
Preemption notes	9-32
Attributes for dynamic tunnel configuration	9-33
Procedures	9-40
List of procedures	
9-1	Configuring dynamic ingress TE tunnels 9-42
9-2	Configuring dynamic ingress unidirectional TP tunnels 9-43
9-3	Configuring a dynamic ingress IP/MPLS tunnel with FRR signaling 9-45
9-4	Configuring resource affinities 9-48
9-5	Configuring SRLG 9-49
9-6	Creating a dynamic co-routed bidirectional tunnel 9-50
9-7	Modifying a dynamic co-routed bidirectional tunnel 9-55
9-8	Linking a dynamic co-routed bidirectional tunnel to performance monitoring 9-59
9-9	Configuring DiffServ-TE 9-61
9-10	Configuring BFD linkage to IS-IS 9-71
9-11	Displaying dynamic tunnel information 9-72
9-12	Displaying auto-sizing statistics 9-77
9-13	Configuring RSVP-TE graceful restart helper mode 9-79
9-14	Configuring T-LDP graceful restart helper mode 9-80
9-15	Configuring MPLS fast reroute 9-82

9-16	Returning MPLS fast reroute to default values	9-83
9-17	Creating an auto facility bypass profile	9-84
9-18	Configuring an auto facility bypass profile	9-87
9-19	Returning an auto facility bypass profile to default values	9-88
9-20	Deleting an auto facility bypass profile	9-89
9-21	Creating a fast reroute profile	9-90
9-22	Configuring a fast reroute profile	9-92
9-23	Returning the fast reroute profile to default values	9-93
9-24	Deleting a fast reroute profile	9-94
9-25	Creating a fast reroute tunnel	9-95
9-26	Creating a facility bypass tunnel	9-96
9-27	Configuring a fast reroute tunnel	9-98
9-28	Returning fast reroute tunnel attributes to default values	9-99
9-29	Deleting a fast reroute tunnel	9-100
9-30	Configuring the IP interface with TE attributes	9-101
9-31	Returning the IP interface TE attributes to default values	9-102
9-32	Displaying MPLS fast reroute values	9-103
9-33	Displaying auto facility bypass profiles	9-104
9-34	Displaying fast reroute profiles	9-106
9-35	Displaying fast reroute protected tunnels at ingress	9-108
9-36	Displaying fast reroute protected tunnels at transit	9-110
9-37	Displaying IP interface information with TE attributes	9-111
9-38	Configuring the soft preemption timer on a node	9-112
9-39	Creating an MPLS-TE unidirectional dynamic LSP with soft preemption	9-113
9-40	Displaying advertised bandwidth for a link	9-114

Additional tunnel configuration

10-1

Ingress/egress CoS mapping 10-1

List of procedures

10-1	Configuring CoS profiles for MPLS tunnels	10-7
10-2	Displaying a summary of CoS profiles for MPLS tunnels	10-9
10-3	Moving co-routed tunnels to a new path	10-10
10-4	Configuring a static co-routed primary tunnel with a dynamic associated tunnel as backup	10-15
10-5	Configuring a dynamic associated tunnel as primary and a static co-routed tunnel as backup	10-17
10-6	Configuring a primary tunnel with an explicit path and a backup tunnel without an explicit path	10-19
10-7	Configuring a primary tunnel without an explicit path and a backup tunnel with an explicit path	10-21
10-8	Configuring a primary tunnel and a backup tunnel without explicit paths	10-23

LDP configuration **11-1**

List of procedures

- 11-1 Configuring LDP globally 11-2
 - 11-2 Displaying LDP configuration 11-3
 - 11-3 Configuring LDP authentication 11-4
 - 11-4 LDP LSP configuration 11-5
 - 11-5 Creating an MPLS TP-TE gateway 11-7
-

Virtual circuit configuration **12-1**

- Changing the PW control word 12-3
- PW group, infrastructure PW, and PW group switchover support 12-3
- Logging for PW state change/up time 12-4
- Virtual circuit statistics at the S-PE NE node 12-4

List of procedures

- 12-1 Configuring dynamic virtual circuits 12-6
 - 12-2 Modifying dynamic virtual circuits 12-9
 - 12-3 Configuring static virtual circuits 12-10
 - 12-4 Modifying static virtual circuits 12-14
 - 12-5 Displaying virtual circuits 12-15
 - 12-6 Configuring L2-VPN virtual circuit CoS profile 12-23
 - 12-7 Displaying L2-VPN virtual circuit CoS profiles 12-24
 - 12-8 Configuring virtual circuit VCCV profiles 12-25
 - 12-9 Displaying virtual circuit VCCV profiles 12-28
 - 12-10 Configuring MS-PW stitching 12-29
 - 12-11 Configuring PW-group protection example 12-32
-

Virtual switch configuration **13-1**

List of procedures

- 13-1 Configuring Control Frame Tunneling over MPLS 13-3
- 13-2 Attaching virtual circuits to a virtual switch multi-segment pseudowire 13-4
- 13-3 Creating an MPLS management virtual switch 13-5
- 13-4 Changing the management virtual switch 13-7
- 13-5 Configuring multiple VLANs to the MPLS virtual switch 13-8
- 13-6 Creating a VS configuration with UNI only with bundled CVIDs 13-11
- 13-7 Configuring ingress and egress I2-transforms 13-13
- 13-8 Configuring the encapsulation egress I2-transform CoS values 13-17
- 13-9 Configuring untagged frames 13-18
- 13-10 Configuring egress I2-transform CoS values 13-20
- 13-11 Configuring CoS profiles for MPLS pseudowire 13-21
- 13-12 Configuring MEF8 TDM services over MPLS 13-22
- 13-13 Displaying CoS profiles for MPLS tunnels 13-24
- 13-14 Displaying CoS profiles 13-25
- 13-15 Displaying virtual switch information 13-27
- 13-16 Displaying MAC table information 13-29
- 13-17 Configuring dual-tag EVPL classification 13-30

L2 VPN service configuration **14-1**

List of procedures

- 14-1 Configuring a 39XX/51XX LSR 14-2
 - 14-2 Configuring a 39XX/51XX VPWS 14-5
 - 14-3 Configuring a 39XX/51XX VPLS 14-9
 - 14-4 Configuring a 39XX/51XX H-VPLS 14-13
-

Interface configuration **15-1**

List of procedures

- 15-1 Creating an MPLS remote management interface 15-2
 - 15-2 Displaying remote interface configuration 15-4
 - 15-3 Displaying interface information 15-5
 - 15-4 Deleting a management virtual switch 15-7
 - 15-5 Adding a static ARP entry 15-8
 - 15-6 Displaying static ARP entries 15-9
 - 15-7 Allocating resources for an MPLS management virtual switch (3916, 3926m, 3928, 3930 and 3931 platforms) 15-10
-

MPLS troubleshooting **16-1**

List of procedures

- 16-1 Displaying FIB information 16-2
- 16-2 Displaying the AIB table 16-3
- 16-3 Displaying FIB entries 16-4
- 16-4 Clearing all FIB or AIB entries 16-6
- 16-5 Enabling logging of FIB or AIB events 16-7
- 16-6 Modifying the AIB timeout for MPLS adjacency 16-8
- 16-7 Running a traceroute on a tunnel 16-9
- 16-8 Running a traceroute on a virtual circuit 16-14
- 16-9 Running a VCCV traceroute 16-17
- 16-10 Running ping on a tunnel 16-20
- 16-11 Running ping on a virtual circuit 16-26
- 16-12 Running VCCV ping for MPLS virtual circuits 16-28
- 16-13 Switching over to the backup TE tunnel 16-31
- 16-14 Switching over to the backup TP tunnel 16-32
- 16-15 Manually forcing a switchover to the standby pseudowire and reverting it back 16-33
- 16-16 Clearing MPLS transit tunnel information 16-34
- 16-17 Clearing GMPLS TP tunnel information 16-35
- 16-18 Clearing traffic statistics for a virtual circuit 16-36

Bidirectional Forwarding Detection configuration **17-1**

Procedural flow for BFD dampening 17-1

List of procedures

- 17-1 Enabling or disabling BFD 17-8
- 17-2 Managing BFD profiles 17-10
- 17-3 Managing BFD HW acceleration 17-13
- 17-4 Setting BFD HW acceleration to default values 17-14
- 17-5 Managing VCCV BFD HW acceleration 17-15
- 17-6 Setting VCCV BFD HW acceleration to default values 17-16
- 17-7 Managing BFD sessions 17-17
- 17-8 Displaying BFD session information 17-20
- 17-9 Displaying BFD default profiles 17-22
- 17-10 Single hop IP-BFD example 17-24
- 17-11 Single hop IP-BFD default profile example 17-26
- 17-12 Multi-hop IP-BFD default profile example 17-27
- 17-13 BFD MPLS-BFD configuration example 17-28
- 17-14 BFD over static MPLS tunnels example 17-29
- 17-15 Dynamic MPLS-TE tunnels example 17-30
- 17-16 Static co-routed MPLS-TP tunnels example 17-31
- 17-17 Static associated MPLS-TP tunnels example 17-32
- 17-18 Dynamic associated MPLS-TP tunnels example 17-33
- 17-19 Static co-routed primary tunnel and dynamic associated backup tunnel example 17-34
- 17-20 Dynamic associated as primary tunnel and static co-routed as backup tunnel example 17-35
- 17-21 Static co-routed primary tunnel and static associated backup tunnel example 17-36
- 17-22 Static associated primary tunnel and static co-routed backup tunnel example 17-37
- 17-23 VCCV BFD over static SS-PW over single LSP example 17-38
- 17-24 VCCV BFD over static MS-PW over single LSP example 17-41
- 17-25 BFD dampening on BFD over static co-routed LSP example 17-45
- 17-26 Displaying BFD dampening profiles 17-48
- 17-27 Displaying the default BFD dampening profile 17-49
- 17-28 Managing BFD dampening profiles 17-50

Alarm Indication Signal with Link Down Indication **18-1**

List of procedures

- 18-1 Displaying AIS/LDI global state 18-2
- 18-2 Enabling and disabling AIS/LDI global state 18-3
- 18-3 Displaying AIS/LDI profiles 18-4
- 18-4 Configuring AIS/LDI profiles 18-5
- 18-5 Configuring AIS/LDI on a static ingress associated TP tunnel (on an LER) 18-9
- 18-6 Configuring AIS/LDI on a static transit unidirectional TP tunnel (on an LSR) 18-11

- 18-7 Configuring AIS/LDI on a static ingress co-routed TP tunnel (on an LER) 18-12
- 18-8 Configuring AIS/LDI on a static egress co-routed TP tunnel (on an LER) 18-13
- 18-9 Configuring AIS/LDI on a static transit co-routed TP tunnel (on an LSR) 18-14
- 18-10 Displaying AIS/LDI global statistics 18-16
- 18-11 Clearing AIS/LDI global statistics 18-18
- 18-12 Example topology: AIS over static MPLS-TP associated tunnels 18-19
- 18-13 Example topology: AIS over static MPLS-TP co-routed tunnels 18-22

Virtual private LAN service integrated routing and bridging 19-1

- VPLS IRB virtual switch configuration 19-1
- VPLS IRB IP interface 19-2
- Unicast IPv4 hardware forwarding over a VPLS IRB IP interface 19-4
- IPv4 local termination and origination over a VPLS IRB IP interface 19-4
- Broadcast/multicast L2 local termination and origination over a VPLS IRB IP interface 19-5
- VPLS IRB IP interface supported functionality 19-5
- IRB specific IFP resources management 19-5
- Scalability 19-6
- Feature considerations 19-6

List of procedures

- 19-1 Reserving the resource manager advanced-l3 feature resource for the 3930, 3932 and the 5160 19-10
- 19-2 Configuring the port to the list of system reserved ports for VPLS IRB 19-11
- 19-3 Enabling or disabling VPLS IRB 19-12
- 19-4 Clearing VPLS IRB statistics 19-13
- 19-5 Configuring an MPLS VPLS IRB virtual switch 19-14
- 19-6 Configuring an MPLS VPLS IP interface associated with the VPLS IRB virtual switch 19-15
- 19-7 Configuring static ARP over the virtual IP interface 19-16
- 19-8 Configuring OSPF passive interface on the MTU for IP interfaces connecting towards eNBs 19-17
- 19-9 Configuring OSPF area route filtering to separate OSPF route distribution 19-18
- 19-10 Displaying system reserved ports for VPLS IRB 19-19
- 19-11 Displaying the status of MPLS VPLS IRB 19-20
- Example 19-20
- 19-12 Displaying VPLS IRB bandwidth usage statistics 19-21
- 19-13 Displaying an MPLS VPLS IRB virtual switch 19-22
- 19-14 Displaying MPLS VPLS IRB IP interface information 19-23
- 19-15 Displaying the static ARP entry over the VPLS IRB IP interface 19-26
- Example 19-26
- 19-16 Configuration example using integrated routing and bridging and MTUs 19-27
- 19-17 Configuration example using MTUs alternative using IP interface 19-30
- 19-18 Configuration example using aggregate gateway (5160) 19-34
- 19-19 Configuration example using aggregate gateway (5160) 19-37

Seamless MPLS configuration 20-1

- IP prefix lists 20-6
 - Autonomous system path access lists 20-7
 - Community lists and expanded community lists 20-7
 - BGP route map 20-8
 - BGP Outbound Route Filtering (ORF) capability 20-10
 - ORF usage in the Ciena solution 20-10
 - ORF generic CLI samples 20-11
 - BGP fast reroute's role in the BGP PIC feature 20-12
 - BGP OAM 20-12
 - BGP-LU solution (AN to AN) 20-13
 - BGP-LU solution (AGN to AN) 20-16
 - BGP-LU use cases 20-17
 - IP over BGP LU features 20-19
 - Management and control traffic 20-19
 - AGN IP active routes selection 20-21
 - Various IP traffic forwarding paths in an AGN 20-22
 - IP High Availability (HA) 20-22
 - IP over MPLS on AN node 20-22
 - Supported activities 20-23
 - Frame receipt on AN 20-24
 - Pseudowire commands 20-26
 - Label ranges 20-27
 - Procedure 20-27
 - 20-1 Sample BGP configuration for a three-node topology with PIC 20-28
-

Flow loop detection for MPLS services 21-1

- Loop detection 21-2
 - Loop detection through MAC motion support for the MPLS network 21-4
- Loop notification 21-5
- Loop prevention 21-5
- Configuration examples for MPLS services 21-5
 - Loop detection through MAC motion on an MPLS network—Example 1 21-6
 - Loop detection through MAC motion on an MPLS network—Example 2 21-6
 - Loop detection through MAC motion on an MPLS network—Example 3 21-7
 - Loop detection through MAC motion on an MPLS network—Example 4 21-7
 - Loop detection through MAC motion on an MPLS network—Example 5 21-8

List of procedures

- 21-1 Enabling and disabling flow loop detection 21-9
- 21-2 Enabling MAC learning at a virtual switch 21-10
- 21-3 Configuring flow loop detection 21-11
- 21-4 Resetting flow loop detection to default values 21-13
- 21-5 Clearing loop detection statistics on the specified port(s) 21-14
- 21-6 Displaying flow learning information 21-15
- 21-7 Displaying flow loop detection information 21-16

Loop-free alternate fast reroute configuration **22-1**

Loop-free alternate 22-2
 Remote loop-free alternate 22-5
 Reversion 22-8
 Micro-loop prevention 22-8
 Reserved IP address for loop-free alternate 22-9
 Configuration and operations 22-10
 Hardware support 22-11
 Notes 22-11

List of procedures

22-1 Enabling or disabling global LFA 22-14
 22-2 Setting or unsetting LFA on a protocol or level 22-15
 22-3 Enabling or disabling LFA on an IP interface 22-16
 22-4 Setting or unsetting LFA on a route map 22-17
 22-5 Setting or unsetting the LFA calculation delay timer after IGP convergence 22-18
 22-6 Setting or unsetting the reserved IP address for LFA 22-19
 22-7 Setting the local micro-loop prevention delay timer for IS-IS 22-20
 22-8 Setting or unsetting the network type of the ISIS protocol on an IP interface 22-21
 22-9 Setting or unsetting the IP interface as a point-to-point interface 22-22
 22-10 Managing route maps 22-23
 22-11 Managing prefix lists 22-24
 22-12 Displaying LFA information 22-25

Sample MPLS topology **23-1**

List of procedures

23-1 H-VPLS configuration example mixed platform 23-2
 23-2 VPLS with CFM configuration example 3916 and 3930 23-9
 23-3 G.8032 and VPLS interoperability example 23-13
 23-4 MPLS-TP configuration example 23-18
 23-5 End-to-End EPL service over MPLS example 23-20
 23-6 Resource affinity configuration for a TE interface example 23-23
 23-7 Encapsulation configuration examples 23-25

About this document

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

The following table lists the version of SAOS software required to run MPLS.

Platform	Requirements
3916	Requires 6.9.0 or later. Supported on all ports.
3926m	Requires 6.17.0 or later. Supported on all ports.
3928	Requires 6.17.0 or later. Supported on all ports.
3930	Requires 6.9.0 or later. Supported on all ports.
3931	Requires 6.9.0 or later. Supported on all ports.
3932	Requires 6.11 or later. Supported on all ports.
3942	Requires 6.12.0.x or later. Supported on all ports.
5142	Requires 6.11 or later. Supported on all ports.
5150	Requires 6.9.0 or later. Supported on all ports.
5160	Requires 6.11 or later. Supported on all ports.

Conventions used in this document

Hyperlinks are indicated by [blue](#) text in this document.

In procedures, the following text conventions are used:

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

Command syntax

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

Symbol	Description
< >	Encloses a variable or literal value that must be specified. Some examples include: server <IpAddress> priority <NUMBER: 1-7> dns <on off> description <String[31]> For server <IpAddress>, the attribute could be entered as server 10.10.11.100 or server www.ciena.com. With priority <NUMBER: 1-7> the text within <> indicates that 1 - 7 are valid values. In the example of dns <on off>, either the literal value of on or off is valid, such as dns on. For description <String[31]>, any string of up to 31 characters is entered.
{ }	Encloses a required value or list of required arguments. One or more values or arguments can be specified. For example, in the syntax: cfm mip create {vlan <VlanId>} {port <PortNameList>} [level <NUMBER: 0-7>] The vlan and port arguments are required. The level argument is optional.
	Separates mutually exclusive items in a list, only one of which can be entered. For example, in the syntax: dhcp client options set subnet <on off> Either on or off must be specified, for example: dhcp client options set subnet on
[]	Encloses an optional value or a list of optional arguments. One or more values or arguments can be specified. For example, in the syntax: arp show [interface <Interface>] You can enter a value for interface <Interface> or not. For example: arp show

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 Service Delivery, Aggregation and Virtualization Switches documentation suite, see *39XX/51XX Service Delivery, Aggregation and Virtualization Switches Product Fundamentals*.

New in this release

The following section summarizes documentation changes in *39XX/51XX Service Delivery, Aggregation and Virtualization Switches MPLS Configuration* for SAOS 6.18.1.

Changes in Revision A

MPLS Dual-homing with Y-cable. Feature JE-73805. This impacts the chapter [“MPLS configuration fundamentals”](#). We made the following changes related to this feature.

- We added the following sections:
 - [“MAC withdraw triggered by inverse VLLI in a Y-cable topology”](#).
 - [“MAC withdraw triggered by inverse VLLI in the active/active model”](#).
 - [“MAC withdraw triggered by inverse VLLI in the active/standby model”](#).

3916: VCCV BFD (CC-1) with Infrastructure PW. Feature JE-78640. This feature impacts the following chapters: [“MPLS configuration fundamentals”](#), [“Virtual circuit configuration”](#), and [“Bidirectional Forwarding Detection configuration”](#). We made the following changes related to this feature.

- In the chapter [“MPLS configuration fundamentals”](#), we added to the section [“Bidirectional Forwarding Detection”](#).
- In the chapter [“Virtual circuit configuration”](#), we added to the section [“PW group, infrastructure PW, and PW group switchover support”](#).
- In the chapter [“Virtual circuit configuration”](#), we added the procedure [“Configuring PW-group protection example”](#).
- In the chapter [“Bidirectional Forwarding Detection configuration”](#), we updated the following procedures:
 - [Procedure 17-23, “VCCV BFD over static SS-PW over single LSP example”](#).
 - [Procedure 17-24, “VCCV BFD over static MS-PW over single LSP example”](#).

VCCV BFD on 3916, 3930, 3931. Feature JE-80370. This feature impacts the following chapters: “[MPLS configuration fundamentals](#)” and “[Bidirectional Forwarding Detection configuration](#)”. We made the following change related to this feature.

- In the chapter “[MPLS configuration fundamentals](#)”, we added to the section “[Bidirectional Forwarding Detection](#)”.

Support BFD Dampening on LSP. Feature JE-80371. This feature impacts the following chapters: “[MPLS configuration fundamentals](#)” and “[Bidirectional Forwarding Detection configuration](#)”. We made the following changes related to this feature.

- In the chapter “[MPLS configuration fundamentals](#)”, we added the section “[BFD dampening](#)”.
- In the chapter “[Bidirectional Forwarding Detection configuration](#)”, we added the section “[Procedural flow for BFD dampening](#)”.
- In the chapter “[Bidirectional Forwarding Detection configuration](#)”, we added mention of the default BFD dampening profile in the procedure “[Displaying BFD default profiles](#)”
- In the chapter “[Bidirectional Forwarding Detection configuration](#)”, we added procedures:
 - “[BFD dampening on BFD over static co-routed LSP example](#)”
 - “[Displaying BFD dampening profiles](#)”
 - “[Displaying the default BFD dampening profile](#)”
 - “[Managing BFD dampening profiles](#)”

Primary and backup pseudo wires (PWs) should not be allowed to attach different virtual switches. Feature JE-80374 and enhancement JE-80077. The feature and enhancement impact the chapter “[Virtual switch configuration](#)”. We made the following change related to the feature and enhancement.

- We added a new paragraph at the beginning of the chapter.

Support current stats (Rx/Tx) on static co-routed tunnels at LER. Feature JE-80460. This feature impacts the chapter “[Static tunnel configuration](#)”. We made the following changes related to this feature.

- We added new steps to the procedure “[Displaying TP tunnel information](#)” to display statistics for static ingress or egress TP co-routed tunnels.

LFA (Loop-Free Alternate) FRR. Feature JE-80034. This feature impacts the chapter “[Loop-free alternate fast reroute configuration](#)”. We made the following change related to this feature.

- We added the chapter “[Loop-free alternate fast reroute configuration](#)”.

Throttling of BFD fault notifications to TM/L2VPN. Enhancement JE-79512. This enhancement impacts the following chapters: [“MPLS configuration fundamentals”](#) and [“Bidirectional Forwarding Detection configuration”](#). This enhancement is covered by the feature Support BFD Dampening on LSP.

PM BulkStats feature sends the wrong IP Address in filename or within file when working with MPLS on SAOS devices. Enhancement JE-78033. This enhancement impacts the following chapters: [“Static tunnel configuration”](#) and [“MPLS troubleshooting”](#). We made the following changes related to this enhancement.

- We updated procedures:
 - [Procedure 8-8, “Displaying TP tunnel information”](#)
 - [Procedure 16-17, “Clearing GMPLS TP tunnel information”](#)

Supported features. Fix JE-80713. This fix impacts the chapter [“MPLS configuration fundamentals”](#). We made the following change related to this fix.

- We added the section [“Features not available on 393x and 3916 platforms”](#).

MPLS configuration fundamentals

This chapter provides an overview of the components and protocols required for configuring Multiprotocol Label Switching (MPLS) on the 39XX/51XX Service Delivery, Aggregation and Virtualization Switches.

With MPLS and related protocols, service providers can provide customers with the perspective of a direct connection to a private line or LAN between their sites with enhanced performance, topology discovery to ensure the most efficient route, and interoperability with popular MPLS switching platforms.

MPLS is explained in the following sections:

- [“MPLS layer 2 service architecture” on page 2-3](#)
- [“Routing protocols” on page 2-12](#)
- [“MPLS label operations” on page 2-21](#)
- [“Encapsulation” on page 2-28](#)
- [“Transport tunnels” on page 2-30](#)
- [“Traffic engineering” on page 2-35](#)
- [“Resource Reservation Protocol - Traffic Engineering” on page 2-35](#)
- [“Virtual circuits” on page 2-36](#)
- [“Control word for pseudowire” on page 2-46](#)
- [“Troubleshooting” on page 2-52](#)
- [“Remote management for MPLS” on page 2-54](#)

Fault management for MPLS is explained in the following sections:

- [“Bidirectional Forwarding Detection” on page 2-55](#)
- [“Alarm Indication Signal with Link Down Indication” on page 2-73](#)

MPLS VPLS integrated routing and bridging is explained in the following section:

- [“Virtual private LAN service integrated routing and bridging” on page 19-1](#)

Seamless MPLS configuration is explained in the following section:

- [“Seamless MPLS configuration” on page 20-1](#)

Flow loop detection for MPLS services is explained in the following section:

- [“Flow loop detection for MPLS services” on page 21-1](#)

MPLS supports circuit (voice), cell-switched (business) and packet switching in a single packet-switched network to save capital and operational expenditures. MPLS speeds up traditional IP forwarding with a simplified header and lookup process. MPLS can be used to build seamless and consistent Layer 2 (L2) Virtual Private Networks (VPNs) between two or more remote sites traversing many different physical networks. Ciena supports

- point-to-point L2VPNs (VPWS)
- multipoint L2VPN (VPLS)

A base MPLS network (or cloud) is defined by a region of two MPLS-enabled switches or routers called Label Edge Routers (LERs), one for ingress and one for egress. A set of Label Switch Routers (LSRs) reside within the cloud between the LERs. The connection from the ingress LER to the intermediate LSR(s) to the egress LERs is called the Label Switched Path (LSP). [Figure 2-1](#) shows a sample MPLS network.

Figure 2-1
MPLS network




Table 2-1 summarizes general MPLS terms.

Table 2-1
General MPLS terms

Term	Definition
LSP	Labeled Switched Path. The specific path through a network that a datagram follows based on its MPLS labels.
LSR	Label Switch Router. A device that switches the labels used to route packets. When an LSR receives a packet, it uses the label included in the packet header as an index to determine the next hop on the Label Switched Path (LSP) and a corresponding label for the packet from a look-up table. The old label is then removed from the header and replaced with the new label before the packet is routed forward.
LER	Label Edge Router. This device is on the edge of the MPLS cloud and is responsible for initiating or terminating the LSPs. LERs can be referred to as edge LSRs.

MPLS can be configured in two implementations:

- MPLS-Traffic Engineering (MPLS-TE), also known as IP/MPLS, which provides
 - better usage of network links and resources
 - bandwidth and QoS guarantees
 - resiliency with fast reroute and recovery
- MPLS-transport profile (MPLS-TP), which provides the same QoS, protection and restoration, and OAM inherent in SONET/SDH, in a way that has a familiar look and feel for network operators. It removes MPLS features that are not relevant to transport networks and adds mechanisms that provide support for critical transport functionality such as:
 - connection verification
 - fault monitoring
 - in-band control and management


MPLS layer 2 service architecture

Layer 2 (L2) VPNs are a type of VPN that use MPLS labels to transport data. MPLS is used to engineer MPLS L2 solutions and to facilitate QoS across the network. L 2 VPN service is provided using MPLS pseudowire technology. Two services are enabled using L2VPN MPLS:

- point-to-point L2VPNs, that is, Virtual Private Wire Services (VPWS)
- multipoint L2VPN, that is, Virtual Private Local Area Network Service (VPLS)

Figure 2-2 shows the MPLS service architecture.

Figure 2-2
MPLS service architecture




MPLS pseudowires can be

- configured to signal using the Targeted Label Distribution Protocol (TLDP) between LERs
- statically-configured

Figure 2-3 shows a manually-configured pseudowire.

Figure 2-3
Manually-configured pseudowire



If using TLDP, an Interior Gateway Protocol (IGP) must be present. An IGP supports routing in an IP network within a single autonomous system (AS).

Figure 2-4 shows a statically-configured pseudowire.

Figure 2-4
Statically-configured pseudowire




Table 2-2 provides definitions of terms used to describe MPLS L2 VPNs.

Table 2-2
MPLS L2 VPN terms


Term	Definition
Provider Edge (PE) Router	Router responsible for encapsulating client data to be carried across the MPLS tunnels. This router is an LER that delivers point-to-point or multipoint L2 connectivity service to the service provider's clients.
Provider router	LSR that operates on the MPLS tunnel label.
Attachment circuit (AC)	Represents the client L2 circuit on the UNI port of the PE.
Pseudowire	A pseudowire can be a single-segment pseudowire (SS-PW) or a MS-PW. An SS-PW consists of a virtual circuit that encapsulates the client payload and adds the pseudowire header. An MS-PW connects pseudowires between different PSN domains which may or may not use the same control plane protocol to control their pseudowires.
Multi-tenant Unit Switch (MTU-s)	A device that support L2 switching functionality and does all the normal bridging functions of learning and replication on all of its ports, including the spoke virtual circuit port. Only a single pseudowire is required per VPLS instance, reducing the signaling overhead between the MTU-s and the PE.

VPWS

VPWS provides point-to-point connectivity between two remote Local Area Networks (LANs). With this type of connectivity, MPLS provides the equivalent of an E-LINE service (EPL or EVPL). Traffic is carried from a single attachment circuit to exactly one pseudowire, and MAC addresses are not learned for the Ethernet attachment circuit.

Figure 2-5 provides an example of a VPWS with a one-to-one connection between CE1 and PE1.


Figure 2-5
VPWS connectivity



VPLS

VPLS provides point-to-multipoint inter-LAN connectivity. With VPLS, PE routers are connected to each other with a full mesh of virtual circuits. Figure 2-6 shows a VPLS full mesh.

Figure 2-6
VPLS full mesh



Each PE provides one or more connections to the customer edge (CE) devices by means of an attachment circuit to the PE MAC/VLAN bridge function. PEs are connected to each other using virtual circuits. The bridge function of the PE attaches to an emulated LAN, which is defined by a virtual switch. The emulated LAN comprised of “VPLS Forwarders” which connect to virtual circuits. This set of VPLS Forwarders connected by virtual circuits represents a VPLS instance.


The bridge function learns MAC addresses, associates them with virtual circuits, and ages them out in a standard manner. When a broadcast/multicast/unknown frame arrives at the bridge function, the frame is forwarded over all the virtual circuits attached to the VPLS forwarder. When a frame arrives on a virtual circuit, the bridge function performs normal Source Address (SA) MAC learning to associate the MAC address with the virtual circuit. In this way, the PE emulates the behavior of a normal LAN. The CE devices appear to be connected to a LAN even though the underlying infrastructure is an MPLS network.

Hierarchical VPLS

The deployment of large-scale VPLS networks where each PE is connected to all other PEs using a full mesh of virtual circuits does not allow complete scalability. As a result, a Hierarchical VPLS (H-VPLS) model is used to allow spoke connections to the VPLS core.

As shown in [Figure 2-7](#), a device provides the functionality to interface with the VPLS core by functioning as an MTU-s, which is connected as a spoke in the VPLS core using a virtual circuit.

Figure 2-7
H-VPLS with MTU-s spoke



In this H-VPLS model, there is only one logical connection, that is, virtual circuit, from an MTU-s to the PE for a given VPLS instance. Each VPLS instance supported by an MTU-s has a virtual switch defined as a virtual L2 switch.

The MTU-s must map incoming frames on its bridge module into VPLS instances. This can be based on one of the following methods:

- physical ports
- VLAN tag of the ingress frame
- virtual circuit MPLS label of the ingress frame
- untagged frames

VPLS membership and MAC learning

The virtual switch provides a mapping between customer-facing attachment circuits and network-facing virtual circuits. A virtual switch can be dual-homed such that there are multiple virtual circuits connecting the VPLS instance to the VPLS hub. MAC addresses to be included in forwarding for the VPLS instance are learned in one of two ways:

- unqualified learning
- qualified learning

MAC addresses that have been learned dynamically are also unlearned or removed explicitly through MAC withdraw mechanisms.

In unqualified learning, all traffic from a specific bridge port is assigned to a single VPLS instance, that is, per-port attachment circuit, and shares a single broadcast domain. MAC addresses need to be unique and non-overlapping among customer VLANs or else they cannot be differentiated within the VPLS instance.

An example of this is port-based VPLS service for a given customer. In this case, all traffic that comes in on a physical port, which can include multiple customer VLANs, is mapped to a single VPLS instance.

In qualified learning, each customer VLAN is assigned to its own VPLS instance, which means that each customer VLAN has its own broadcast domain and MAC address space. Unlike unqualified learning, MAC addresses among customer VLANs can overlap with each other. They are handled correctly since each customer VLAN has its own Forwarding Information Base (FIB). When a VPLS instance is defined per-port per-VLAN, the customer VLAN must be the same on each bridge port that is joining the virtual switch. An example of qualified learning is per-port per-VLAN VPLS service.

MAC withdraw

MAC withdraw is a signal typically sent from a node that wants remote peers to flush all learned MAC addresses in a given VPLS instance. The receiving node unlearns either MAC addresses present in the signal or all MAC addresses if they are absent in the signal. SAOS only supports empty MAC withdraw (no MAC address in the signal) to indicate to flush all learned MAC addresses for a given VPLS instance.

The send and receive processing of the MAC address withdraw signal capability is, by default, turned on. It cannot be turned off. The MAC withdraw signal is sent only when the standby pseudowire becomes operational. It is not used for any other pseudowires that are not part of a protection group.

MAC withdraw signaling is sent in band with the user data traffic over the pseudowire OAM channel 0x04 GAL/G-ACh with an ACh type of 0x028.

MAC withdraw triggered by inverse VLLI in a Y-cable topology

If inverse VLLI (iVLLI) is configured in a Y-cable topology, MAC withdraw can be triggered when there is a change in the Y-cable state. The logic and intelligence of when and where the trigger is necessary lies entirely in the iVLLI. To transmit the MAC withdraw message, static PWs use the methods described in this section.

This section describes the methods used to transmit the MAC withdraw messages in the following scenarios:


- active/active PW
- active/standby PW

Note: For information on configuring inverse VLLI using the active/active and active/standby models, see the chapter “Virtual Link Loss Indication configuration” in *39XX/51XX Service Delivery, Aggregation, and Virtualization Switches Advanced Ethernet Configuration*.

MAC withdraw triggered by inverse VLLI in the active/active model

The following figure shows the active/active model.


Figure 2-8
Y-cable active/active model



In this active/active model, both PW-1A and PW-2A are in the active state. There is a one-to-one mapping between the iVLLI instance and the PW. When a PW operational-state change event occurs, iVLLI directs the registered PW in a VPLS instance to issue the MAC withdraw message.

The following figure shows the active/active model, with a Y-cable fault.

Figure 2-9
Y-cable active/active model with a fault and MAC withdraw




If a fault occurs on the client port of device PE-1A, the sequence of events is as follows:

- 1 PE-1A detects a Y-cable fault.
- 2 PE-2A Y-cable comes up because of the iVLLI mechanism.
- 3 PE-2A sends a MAC withdraw signal to device Z.
- 4 Z processes the MAC withdraw signal and flushes the corresponding forwarding table entries.

MAC withdraw triggered by inverse VLLI in the active/standby model

The following figure shows the active/standby model.

Figure 2-10
Y-cable active/standby model




In this active/standby model, devices PE-1A and PE-2A are connected by a mesh PW. This mesh PW is part of the same VPLS instance with the active/standby PW-1A. In this topology the iVLLI instance registers for this mesh PW. When the fault occurs, iVLLI directs the mesh PW to issue the MAC withdraw message.

The following figure shows the active/standby model, with a Y-cable fault.

Figure 2-11

Y-cable active/standby model with a fault and MAC withdraw



If a fault occurs on the client port of device PE-1A, the sequence of events is as follows:

- 1 PE-1A detects a Y-cable fault.
- 2 PE-2A Y-cable comes up because of the iVLLI mechanism.
- 3 PE-2A sends a MAC withdraw signal to device PE-1A over the mesh PW.
- 4 PE-1A processes the MAC withdraw signal and flushes the corresponding forwarding table entries.

Routing protocols

Routing can be static or dynamic.

Static routing

Static routes are not automatically updated when IP network topologies change. Static routes are manually reconfigured to adapt to new network topologies.

The egress interface is an optional parameter that can be specified while configuring the static route over any numbered or unnumbered IP interface. This allows the user to choose the “best” path if there are multiple equal-cost unnumbered paths to the given next-hop while using static routing.

Dynamic routing

With dynamic routing, MPLS network topology and routing information is monitored and maintained by the following routing protocols:

- Open Shortest Path First (OSPF)
- Open Systems Interconnect (OSI) Intermediate System to Intermediate System (IS-IS) Intra-domain Routing Protocol

These routing protocols discover neighbors in the MPLS network and maintain the optimum route to them.

Open Shortest Path First (OSPF)

MPLS network topology and routing information is monitored and maintained by OSPF. OSPF discovers neighbors in the MPLS network and maintains the optimum route to them. OSPF is an IGP,

OSPF is used to create an IP routing table for building dynamic MPLS tunnels. OSPF is a link-state protocol that uses configurable metrics to associate a cost with a link. These metrics allow network administrators to manage their network based on the speed, reliability, and delay of the network.

The OSPF protocol is a link-state routing protocol, which means that the routers exchange topology information with their nearest neighbors. The topology information, in the form of a link-state advertisement (LSA), is flooded throughout the AS, so that every router within the system has a complete representation of the topology. Devices build a Link State Database (LSDB) based on this information. Each Area Border Router (ABR) has one LSDB for each area to which it is connected. This information is then used to calculate the shortest end-to-end paths through the system. This is accomplished by means of Dijkstra's Shortest Path First (SPF) algorithm. The SPF tree, also known as the best path tree, is then submitted to the routing table as OSPF routes. When a network topology change occurs, a recalculation of the shortest path tree is performed. The network will converge when all routers have recalculated their routing tables as a result of a change in the topology.

OSPF neighbors are any two routers that have an interface to the same network. When an OSPF device first joins a network, it uses the Hello Protocol to discover its neighbors. Neighbors may form adjacencies for the purpose of exchanging routing information. Not all neighbor pairs can become adjacent. Adjacencies are formed by synchronizing the neighbors' topology databases through the database exchange process. Two devices are said to be fully adjacent when they have synchronized their databases. Routing information is exchanged between the adjacent routers only, thereby conserving bandwidth. Also, an authentication mechanism prevents unauthorized neighbors from establishing adjacencies.

The multi-level hierarchy (two-level for OSPF) called “area routing” allows the information about the topology within a defined area of the AS to be hidden from routers outside this area, enabling an additional level of routing protection and a reduction in routing protocol traffic. The authentication of all protocol exchanges prevents unauthorized routers from joining the AS. The system software supports two OSPF areas.

Intermediate system to intermediate system protocol (IS-IS)

Intermediate system to intermediate system (IS-IS) is a Layer 2 (L2) link-state interior gateway protocol (IGP) used for intra-domain routing. Each router (IS) maintains a link-state database (LSDB) which contains information about neighbors and interfaces. The LSDB is used by a Shortest Path First (SPF) algorithm to calculate the shortest path to all known destination nodes.

IS-IS can be used to replace OSPF functionality or it can be implemented simultaneously with OSPF.

Protocol Data Units

In IS-IS, there are three types of packets, or Protocol Data Units (PDUs).

Table 2-3
Types of PDUs

Type	Description
IS-IS Hello PDUs (IIHs)	IIHs discover, establish, and maintain IS-IS neighbor adjacencies. Point-to-point (unicast) network types exchange point-to-point IIHs. Broadcast (multicast) network types exchange two types of IIHs: L1-LAN IIHs and L2-LAN IIHs.
Link State PDUs (LSPs)	LSPs share routing information by announcing the state of the links that are attached to them.
Sequence Number PDUs (SNPs)	SNPs are a combination of one or more LSPs. They can be complete SNPs (CSNPs) or partial SNPs (PSNPs). CSNPs send a complete summary of the LSDB for that IS for that level. PSNPs acknowledge and request LSPs and are typically used for updates after a route adjacency change or for sending additional information.

Addressing

Network Entity Titles (NETs) are used in IS-IS to identify routers. NET addresses are variable in length, range from 8 to 20 octets, and are read from right to left.

Using the example NET address 49.0001.1921.6800.1002.00, it is read as:

- 1-byte Selector: 00
- 6-byte System ID: 1921.6800.1002
- 13-byte (variable) Area ID:
- 49.0001

Note: The Area ID is comprised of the 1-byte Authority and Format Indicator (AFI), 49, and the 1-12-byte Domain ID, 0001.

If a packet is destined for a different area, it is routed based on the area address. Packets destinations in the same area are routed based on the System ID and packet destinations in other areas are routed based on the Area ID.


Areas and routing levels

IS-IS routing domains comprise end systems (nodes) and intermediate systems (routers) and can be subdivided into routing subdomains called areas. Each router resides in one area and floods routing information from its LSDB using link-state PDUs (LSPs) to other adjacent routers, which are determined by the routing level assigned to each router. Routers store the information from the LSPs they receive in their LSDBs.

Routers can be configured as:

- L1 intra-area router for routing within the same area
- L2 inter-area router for routing between areas
- L1-L2 router that performs both functions

Figure 2-12
Domain with intra-area and inter-area route exchanges



An L1 router:

- knows only about the topology of its own area, which is learned from LSPs flooded by L1 or L1L2 routers in the same area
- sends L1 LSPs
- forms L1 adjacencies with other L1 and L1L2 routers in the same area
- maintains an L1 LSDB that is identical to other L1 routers in the same area
- is not connected to the backbone and sends inter-area traffic through the closest L2 router

An L2 router:

- is a backbone router
- sends L2 LSPs
- forms an L2 adjacency with L2 and L1L2 routers in all areas
- does not form adjacencies with L1 routers
- maintains an L2 LSDB for other L2 routers in all areas
- can be configured to summarize redistributed routes to the L1 LSDB


Note: Unlike OSPF, no IS-IS area functions solely as the backbone. The backbone is formed by a series of L2 routers from different areas.

An L1-L2 router:

- is a border router that connects intra-area routers with inter-area routers
- sends L1 and L2 LSPs
- sets an attach (ATT) bit in L1 LSPs to inform L1 routers about a default route to it
- forms L1 adjacencies with L1 routers in the same area
- forms L2 adjacencies with L2 routers in all areas
- forms both L1 and L2 adjacencies with L1-L2 routers in the same area
- forms L2 adjacencies with other L1-L2 routers from different areas
- maintains L1 and L2 LSDBs
- can be configured with static routes by means of address prefixes (for faster routing table convergence)
- can be configured to summarize routes from the L1 LSDB and add them to L2 LSPs
- can be configured to leak routes from the L2 LSDB to the L1 LSDB

The following figure shows a sample Level 1 topology:


Figure 2-13
Sample Level 1 topology



The following figure shows a sample Level 2 adjacency with two areas:

Figure 2-14


Sample Level 2 adjacency with two areas



The following figure shows a sample Level 1-2 adjacency with four areas:

Figure 2-15

Sample Level 1-2 adjacency with four areas



Interface types

IS-IS supports point-to-point (unicast) and broadcast (multicast) interface types. The interface types must match on both sides of a link or the adjacency will not establish.

Three-way handshake for the point-to-point interface

The IS-IS feature supports the extension to the protocol, defined by RFC5303, that provides a process for using three-way handshakes to form point-to-point adjacencies.

Route summarization from L1 to L2

Multiple groups of addresses can be summarized from L1 into L2. Routes learned from other routing protocols can also be summarized. The default cost used for summary route is 20. A maximum of 1500 summary routes can be configured.

Route leaking from L2 to L1

Route leaking from L2 to L1 ensures that the L1 router in an area chooses the most optimal path to reach prefixes outside of its own area. When route leak is enabled on the L1-L2 router, it redistributes prefixes from L2 to L1 so that the level 1 routers in the area can choose the most optimal path to reach prefixes outside the area. By having more detail about inter-area routes, an L1 router can make a better choice with regard to which L1-L2 router to forward the packet. To control advertising of routes not generated inside the area, an up/down bit indicates whether the route defined in the TLV has been leaked. If the up/down bit is set to 0 the route was originated within that L1 area. If the up/down bit is not set (it is 0), the route has been redistributed into the area from L2. The up/down bit is used to prevent forwarding loops. An L1-L2 router does not re-advertise into L2 any L1 routes that have the up/down bit set. A maximum of 1500 routes can be configured.

Attach bit

The attach bit is set by L1-L2 routers to notify L1 routers that they can reach the rest of the network, that is, default routing. Inter-area routing is accomplished by directing all traffic to the nearest L1-L2 router that is part of the backbone that spans multiple areas. L1 routers discover L1-L2 routers that are part of the backbone by looking for the attach bit set in the L1 LSPs generated by these backbone L1-L2 routers. The attach bit is only set by L1-L2 routers that have at least one L2 adjacency that spans multiple areas.

Overload bit

When you set the overload bit, SPF is recalculated on all nodes to exclude the node with overload bit set. Excluding a node from an SPF calculation for a certain period of time is useful for maintenance periods.

Default route/ignore attach bit

Ignore attach bit ignores the attach bit and does not install a default route towards the L1-L2 router. The attach bit can be configured. By default, the ignore attach bit setting is disabled. This means that if an attach bit is received, the default route is installed.

IS-IS appends the default route (with the already configured management default route) on receiving the attach bit and handles the protocol configured default route in the following ways:

- Uses “65536” as the default metric for the default management route.
- All L3 protocol’s IS-IS/BGP use their administrative distance as the metric while configuring protocol-driven default routes.
- Management connectivity remains intact even if the protocol-configured default route is removed from the system.

Routing metrics

IS-IS, like any other protocol, uses metrics to determine the cost of the link. The lower the cost, the better the path.

All IS-IS links use a metric of 128 by default. If the same destination is reachable through both Level 1 and Level 2 paths, then the Level 1 path is preferred.

The metric style can be set to narrow, wide, or transition. Narrow uses the old-style TLVs with link values from 1 - 63. Wide uses the new style TLVs with link values from 1-16777215. Transition enables both narrow and wide styles.

Redistribution

Static route redistribution for L1 and L2 routes is supported. No support is currently provided for other route redistribution.

IS-IS interface MTU enhancement/limitation

The IS-IS interface MTU is derived from an IP interface MTU setting. This value is used by IS-IS to generate hello packets (while hello padding-on) to support MTU discovery. If the user sets the IP interface MTU to greater than 1500, then IS-IS hello packets are treated as jumbo frames. These packets have type 0x8870 (Ether type) as explained in this [draft](#).

As of SAOS 6.18, the IS-IS interface MTU has been detached from the IP interface MTU. It has a fixed value of 1500. There may be interoperability issues with other vendors if they run IS-IS (with hello padding-on) with an interface MTU more than 1500.

LSP zero-age lifetime

Each LSP in the LSDB has a MaxAge. After the MaxAge of an LSP is reduced to zero, the LSP is removed from the LSDB if no LSP with the larger sequence number and the same LSP ID is received to update the local LSP. After the old LSP is removed from the LSDB, it is kept for the period of ZeroAgeLifetime. Zero-age lifetime is the timer that keeps expired prefixes. (In the period of ZeroAgeLifetime, only the header of the old LSP is kept; if the authentication field exists in the old LSP, the authentication field is also kept in this period.)

When ZeroAgeLifetime is expired, the LSP is really removed from the LSDB. In general, the default value of MaxAge is 1200 seconds and the default, non-configurable value of ZeroAgeLifetime is 60 seconds.

Static route preference and switching

The order of preference between L1 and L2 routes (as stated in RFC2966) is as follows:

- L1 route with internal metric (highest preference)
- L2 route with internal metric
- L1 route with internal metric and up/down bit set
- L1 route with external metric
- L2 route with external metric
- L1 route with external metric and up/down bit set (lowest preference)

In case multiple static routes are configured for the same destination with different metrics, then the one with the least metric is chosen. If the chosen least metric route goes down, then there is no fallback to the alternative route. Static routes are configured by the administrator so the system does not monitor them, and hence they are not removed or added dynamically. To resolve this problem, Ciena recommends to always configure all the static routes with the default or equal metric values.

MPLS label operations

A packet enters an MPLS network through the ingress LER. The ingress LER performs a routing lookup to determine the MPLS label, which corresponds to a destination as stored in the forwarding information base (FIB). This label also identifies the Forwarding Equivalence Class (FEC) so that the packet is included in a set of packets with the same forwarding destination and Class of Service (CoS).

MPLS label operations are push, swap, and pop.


The ingress LER inserts (pushes) this MPLS label (or stack of labels) between the L2 and Layer 3 (L3) headers of the IP packet and changes the EtherType value to indicate that the packet is labeled. Table 2-4 shows the insertion of the MPLS label.

Table 2-4
MPLS label operations

Label operation	Description
Push	The ingress LER inserts (pushes) the MPLS label or stack of labels onto the MPLS packet and changes the EtherType value to indicate that the MPLS packet is labeled.
Swap	The ingress LER initiates the LSP and forwards the packet over the LSP to the adjacent LSR. Each LSR swaps the label to the next hop label (or stack of labels) of the next LSR (or egress LER) and forwards the packet.
Pop	When the packet reaches the egress LER, the egress LER performs a routing lookup to determine the egress destination, removes (pops) the MPLS label or labels and delivers the MPLS packet.

Figure 2-16 shows an example of how MPLS labels are used to route packets through the MPLS network. In this example, the ingress LER is upstream to LSR1, and LSR1 is downstream to the ingress LER.

Figure 2-16
MPLS label operations



In the example, route 1.1.1.1 is reachable through LSR1. LSR1 provides label 200 to the Ingress LER, which indicates that traffic can be sent to route 1.1.1.1 by means of label 200. The ingress LER then programs label 200 in its label forwarding information base (FIB). The ingress LER performs a push operation. It assigns label 200 which provides destination and QoS information. LSR1 performs a swap operation: it replaces the ingress LER label with the egress LER label. The egress LER pops label 100: it removes the MPLS label and exposes the client frame.

MPLS labels

An MPLS label is a 32-bit field with the bit values as listed in [Table 2-5](#).

Table 2-5
MPLS label format

Bit name	Label (20 bits)	TC (3 bits)	S (1 bit)	TTL (3 bits)
Bit number	0-19	20-22	23	24-31
Description	20-bit label with bits Note: The values between 0 and 15 are reserved.	3-bit IP precedence CoS	1-bit value indicates whether the label is the last label (single label or bottom of a label stack)	8-bit TTL of IP header prevents infinite loop of the packet


Multiple labels, that is, a stack, are used for the following applications:

- MPLS VPN uses two labels: the top label is for the egress router and the second label is for the VPN
- IP/MPLS uses two or more labels: the top label is the endpoint of the TE tunnel and the remaining labels show the destination

0

In some cases, a combination of three or more labels are used for both. Multiple labels are inserted as shown in [Figure 2-17](#).

Figure 2-17
Example MPLS VPN stack



MPLS labels are:

- implicit NULL label
- explicit NULL label
- router alert label

Implicit NULL label

In the case of Penultimate Hop Pop (PHOP), the penultimate hop LSR pops the top label before forwarding the packet to the egress LSR. The egress LSR signals that the label should be popped at the penultimate hop when it provides an implicit NULL (0x03) label to the preceding hop, for a given tunnel that is terminated by the egress LSR.

Explicit NULL label

The explicit NULL label has label value zero. The implicit forwarding assumptions are the same as for the implicit NULL label in that the egress LSR submits label value zero to its predecessor. The egress LSR signals that at the penultimate hop the label should be swapped with value zero but maintain the 'traffic class (tc) bits' intact so that proper queuing can be applied to the received packet. The label stack depth is maintained but the top label value of zero is only used for QoS purposes.

Router alert label

The router alert label has label value 1. The router alert label is pushed above the pseudowire label and below the tunnel label in the MPLS label stack. When the egress PE router receives the pseudowire packet, the router alert label is exposed and the packet is delivered to the control plane.

Label Distribution Protocol

The principal role of Label Distribution Protocol (LDP) is to establish and maintain virtual circuits based on the agreement of the meaning of the label used to forward traffic. Targeted LDP and Topology LDP sessions are used. Targeted LDP sessions are used to learn labels for dynamic MPLS VC. Topology LDP sessions are used for peer-peer tunnels. Because LDP is a peer-to-peer protocol based on the establishment and maintenance of TCP sessions, the following natural benefits exist:

- LDP messages are reliably delivered by the underlying TCP, and state information associated with explicitly-routed LSPs does not require periodic refresh.
- LDP messages are flow-controlled, that is throttled, through TCP.

The categories of LDP messages are:

- Discovery messages announce and maintain the LSR presence in the network.
- Session messages create, maintain and terminate sessions between two LDP peers that use LDP to exchange label mapping information.
- Advertisement messages create, change and delete label mapping for FEC, which is a set of packets with similar or identical characteristics.
- Notification messages provide advisory information and signal error information.

MPLS TP-TE gateway


The MPLS TP-TE gateway is a gateway between the TP and TE MPLS domains. Connecting the TP and TE MPLS domains provides end-to-end service resiliency.

MPLS TP-TE gateway requires interworking of LDP, T-LDP, and IS-IS with other vendors.

MPLS MP-TE gateway adds a pseudowire mode of gateway. This propagates up to the MTU node and the MTU node switches traffic to the backup PWE and makes the backup virtual circuit active. This way end-end service resiliency is achieved. Only unprotected virtual circuits can be created or added in gateway mode.


The following figure shows the logical model for virtual switch connectivity.

Figure 2-18
Virtual switch connectivity for gateway mode




The following figure shows normal operation across gateway node virtual circuits. The virtual circuit status propagation keeps the end-to-end paths according to the PWE status on each domain. The primary virtual circuit side forwards traffic and the backup virtual circuit blocks traffic.

Figure 2-19
Normal operation across gateway node virtual circuits




Virtual circuit status on MPLS-TP domain are carried as PDUs inside the virtual circuit status PDU messages, as shown in the following figure.

Figure 2-20
Virtual circuit status




The status TLV on the IP/MPLS PWE virtual circuit aligns with T-LDP signaling inside the LDP notification message, as shown below:



When a fault occurs, the gateway node sends the failure notification to the other connected gateway virtual circuit. When the status message reaches the MTU node, the MTU node activates the backup virtual circuit. The backup virtual circuit is activated using the status message. The status message is forwarded as active to the peer backup virtual circuit. In the following figure, the backup virtual circuit is on the IP/MPLS side.

Figure 2-21
End-to-end service resiliency using gateway node virtual circuits



Encapsulation

The packet switching services can be partitioned into different categories based on the packet format, information used to build and process the packet, and the de-multiplexing mechanism used to identify the payload. The pseudowire encapsulations are defined by the IETF standards organization. It is a methodology that specifies how client data from different types of Attachment Circuits (ACs) can be propagated over the PSN.

MEF8 encapsulation

The initial encapsulation supported is MEF8, but it is not fully compliant with MEF8. The CESoETH control word is not supported; instead, the CEP header from RFC 4842 is used. As a result, m-bit events are not supported for PMs. From MEF8, the following are supported:

- The MEF8 header with EC-ID configured as pseudowire channel identifier
- Single or double VLAN tag with configurable VLAN ID, PCP/DEI value per channel, and TPID fixed at 0x8100

While the PW procedures were being defined at the IETF for TDM traffic, the Metro Ethernet Forum adopted an earlier version of the TDM PW standard and retrofitted it to carry TDM as payload on native Ethernet networks. The forwarding semantics of the PW payload is carried in a label that is identical in format to a PW label, but is termed as EC-ID, Ethernet Connection Identifier.

MPLS encapsulation

There are two ways to implement MPLS headers for TDM PWEs:

- Native MPLS or TDMoMPLS
- MPLS Encapsulation or TDMoMEF8oMPLS

Native MPLS supports using an MPLS header in front of the payload where the bottom MPLS label is used as the EC-ID for the far end to demux the connections. As of SAOS 6.17.1, native MPLS is supported for PWEs attached to the PWE module TDM ports and FRU module TDM ports.

MPLS encapsulation uses an MPLS header in front of the MEF8 encapsulated TDM payload. As of SAOS 6.16, MPLS encapsulation of MEF8 encapsulation is supported for PWEs attached to TDM SFP ports or PWE module TDM ports and FRU module TDM ports.

Both dynamic and static MPLS LSPs are supported but only static pseudowire configuration is supported. The supported MPLS configuration is listed in this table:

Table 2-6
Supported encapsulation

Supported encapsulation	MPLS LSP	MPLS VC (PW)
Native MPLS	Static and dynamic	Static only
MEF8	Static and dynamic	Static and dynamic
Dry Martini	N/A	N/A

This table summarizes encapsulation information for TDM services.

Table 2-7
Encapsulation information

Encapsulation type	Description	Supported on
MEF8	The encapsulation header is identified by an Ethertype of 88D8. The encapsulation header comprises: • Emulated Circuit Identifier (ECID), which identifies the emulated circuit being carried • Control Word, which provides sequencing and signaling of defects such as AIS • Optional RTP Header, which provides timing and sequencing by means of Real-Time Transport Protocol (RTP). The RTP header is not supported in SAOS software.	TDM SFP ports 3926m T1/E1 PWE module ports
Dry Martini	The encapsulation header is an MPLS header with an optional control word. The MPLS bottom label is the PWE label. There is no tunnel label.	PWE module ports 3926m T1/E1
MPLS	The encapsulation header is an MPLS header. The MPLS bottom label is the PWE label. There is also a tunnel label identifying the MPLS tunnel. MPLS encapsulation of MEF8 encapsulation is supported for PWEs attached to TDM SFP ports, PWE module TDM ports or 3926m T1/E1.	PWE module ports 3926m T1/E1 TDM SFP ports

Available pseudowire types

Circuit Emulation over PSN (CESoPSN): CESoPSN TDM pseudowire technology supports framed and channelized TDM services over packet switched networks. It uses multiples of the TDM frame itself. CESoPSN packetizes the TDM data.

Structure Agnostic TDM over Packet (SAToP): SAToP (RFC 4553) or Structure Agnostic TDM over Packet, is a TDM pseudowire technology that differs from CESoPSN in that it treats the TDM traffic as a data stream and ignores the framing or the timeslots (DS0). It provides similar functionality in its unframed mode.

For configuration examples, see [“Encapsulation configuration examples” on page 23-25](#).

Transport tunnels

Tunnels can be created dynamically or statically. This table lists the supported tunnels for IP/MPLS and MPLS-TP.

Table 2-8
Supported tunnels for IP/MPLS and MPLS-TP

	IP/MPLS	MPLS-TP
Static	• Unidirectional • Associated bidirectional • Co-routed bidirectional	• Unidirectional • Associated bidirectional • Co-routed bidirectional
Dynamic	• Unidirectional • Associated bidirectional	• Unidirectional • Associated bidirectional • Co-routed bidirectional

MPLS-TP LSPs are established over IP interfaces. RSVP and IGP are enabled over the IP interfaces. The entire MPLS control plane operates over numbered or unnumbered IP interfaces. The IP interfaces for 39XX/51XX switches are configured as logical interfaces over the physical port. For example, the network operator can configure an IP interface over the Ethernet port of a VLAN. More than one IP interface can be configured over a given physical port, which allows for a many-to-one relationship.

Static unidirectional

[Figure 2-22](#) shows an example of a static unidirectional tunnel. The ingress LER is also referred to as the tunnel head. When outside data arriving at the ingress PE needs to be sent to the egress LER over the MPLS infrastructure, a label which is significant to the next router in the tunnel is inserted, that is, pushed, and the encapsulated data is forwarded to the next hop along the

tunnel. The label that was pushed is referred to as the egress label on the router where it is pushed and as the ingress label on the immediate downstream router. This next router then replaces that label with the label understood by the next downstream router, that is, it swaps the label.

Figure 2-22
Sample static unidirectional tunnel




The push-swap operation is repeated until the data arrives at the far-end LER, which is the egress LER. The egress LER removes the label, that is, pops the label. The data is further processed according to the enclosed encapsulation fields.

Transit tunnels are defined on the transit LSRs.

Static associated bidirectional

Figure 2-23 shows a sample associated bidirectional tunnel. In an associated bidirectional tunnel, each end node is configured with an LSP to the other end. In addition, both nodes are configured with the association of the forward and reverse LSP. The association configuration of two counter-directed LSPs at each end node makes the LSP an associated bidirectional tunnel.


Figure 2-23
Sample static associated bidirectional tunnels



Static co-routed bidirectional

Figure 2-24 shows sample co-routed tunnels. In a co-routed bidirectional tunnel, the initiating end node is configured with co-routed bidirectional LSP. The terminating end node creates a reverse LSP upon receiving the signal. Each LSR needs to have an ingress-egress label pair in each direction.

Figure 2-24
Sample static co-routed tunnels



In the case of static co-routed bidirectional tunnels, both end nodes are configured with LSPs in each direction.


Note 1: The term 'ingress' is used interchangeably with the term 'initiator' for bidirectional tunnels to identify the initiator node. The term 'ingress' is preserved as a keyword in the SAOS CLI for backward compatibility.

Note 2: Penultimate Hop Popping (PHP) is prohibited for MPLS-TP tunnels.

Dynamic unidirectional

Figure 2-25 shows an example of a dynamic unidirectional tunnel. A unidirectional tunnel is created using a dynamic label distribution protocol: in the example, RSVP. An RSVP path message is forwarded from the end node at point 1 along the intended path of the tunnel. RSVP Resv messages are distributed starting at point 2a in the exact opposite direction of the tunnel to signal the labels.

Figure 2-25
Sample dynamic unidirectional tunnel




Dynamic associated bidirectional

Figure 2-26 shows a sample dynamic associated bidirectional tunnel. In a dynamic associated bidirectional tunnel, the forward and reverse tunnel must be associated at two ends by user configuration. Each component unidirectional LSP is then dynamically signaled.

Figure 2-26
Sample dynamic associated bidirectional tunnels

Forward dynamic LSP



Dynamic co-routed bidirectional

A dynamic co-routed bidirectional LSP is single sided, that is, it is configured at the initiator end only. The terminating end node creates a reverse LSP dynamically upon receiving the signal. There is no tunnel-specific configuration at the terminating PE.


Dynamic signaling assures LSP resiliency if there are multiple failures in the MPLS network, provided some physical connectivity exists between the source and destination. This is done through periodic, persistent RSVP retries. This eliminates human intervention, which is required for the static configuration.

Note 1: IS-IS IGP is required for dynamic co-routed bidirectional tunnels.

Note 2: Dynamic co-routed bidirectional tunnels are supported for MPLS-TP only.

Figure 2-27 shows a sample dynamic co-routed bidirectional tunnel.

Figure 2-27
Sample dynamic co-routed bidirectional tunnels



Traffic engineering

Traffic engineering (TE) is a method of optimizing the performance of a telecommunications network by dynamically analyzing, predicting and regulating the behavior of data transmitted over that network.

For dynamic co-routed bidirectional tunnels, TE parameters are configured on IP interfaces and LSPs.

For IP interfaces, TE parameters are distributed by the IGP and build the TE topology in each MPLS node. TE parameters are predominantly configured over IP interfaces with an IP interface-to-physical port one-to-one relationship. A consideration for bandwidth-related configurations is that they must share multiple IP interfaces that map to the same physical port.

TE attributes for LSPs provide flexibility in the configuration of LSP attributes for path options. TE preferences for an LSP enable it to serve as transport for the pseudowire service traffic. All the pseudowires mapped on this custom-built LSP transport share the desired Quality of Service when passing through MPLS networks. For instance, service requirements for latency, delay, loss probability and guaranteed bandwidth dictate the attributes associated with the LSP.

For LSPs, TE parameters are used in the CSPF calculation to pick candidate links from the local TE database that meets the LSP TE requirements and then builds the shortest path to the destination.

Resource Reservation Protocol - Traffic Engineering

Dynamic MPLS networks are created and monitored with the signaling of Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for tunnels, and Label Distribution Protocol (LDP) for virtual circuits.

RSVP-TE is a protocol used to establish LSPs in dynamic MPLS networks. LSPs allow the pseudowires to be carried through the network. One LSP can carry several pseudowires. The traffic that flows along an LSP is defined at the ingress node.

To create an LSP tunnel, the first MPLS device on the path creates an RSVP-TE path message. A label binding for this path is requested and indicates which network layer protocol is to be carried over the path. When the sending device finds a path that either meets the tunnel's QoS requirements, satisfies the policies criteria, or can maximize the use of the network resources, an explicit route is specified. This explicit route can be dynamically changed if the device finds a better route. This event is recorded and the sender device is notified.

The LSP tunnel carries the load to the destination device along the path. When the destination is reached, a received message is sent back to the sending device, following the path in reverse order, which establishes a bidirectional LSP.

Virtual circuits

A virtual circuit is a bidirectional connection between endpoints that can multiplex and de-multiplex traffic over tunnels. Multiple virtual circuits between two endpoints can use the same tunnels. For virtual circuits associated with an MTU-s, a secondary virtual circuit can be configured to support dual-homed protection.

Raw and tagged pseudowire type for virtual circuits

The term attachment circuit is used to describe the customer-facing port of an Ethernet L2VPN service. The two types of attachment circuits available are:

- EPL attachment circuit, which is a port added to an L2VPN service
- EVPL attachment circuit, which is a single VLAN added to an L2VPN service

An EPL attachment circuit indicates that the outer VLAN tag of the frame, if present, is not service delimiting and therefore is not meaningful to the PE. An EPL attachment circuit always assumes the outer tag of the frame is a C-Tag.

An EVPL attachment circuit indicates that the outer VLAN tag of the frame is service delimiting and should be used to identify the traffic. An EVPL attachment circuit always assumes the outer tag of the frame is an S-Tag.

Note 1: When an MPLS EPL is added to a UNI port, this UNI port is automatically removed from the default VLAN. For an EVPL, you can remove the UNI Port from the default VLAN by using the CLI. When the default VLAN membership for the EPL/EVPL port has been removed, traffic that is flooded on the default VLAN no longer egresses the UNI port.

Note 2: MPLS virtual circuits share internal resources with MPLS attachment circuits and sub-ports. This means that the scale of these features may be affected by the number of MPLS virtual circuit and attachment circuit instances on the switch.

MPLS forwarding behavior depends on the type of attachment circuit and virtual circuit combination configured. There are four possible combinations:

- EPL attachment circuit with raw virtual circuit
- EVPL attachment circuit with raw virtual circuit
- EPL attachment circuit with tagged virtual circuit
- EVPL attachment circuit with tagged virtual circuit


A raw pseudowire type virtual circuit never carries a service delimiting tag. There are only two possible options for this type: ignore the tag or pop it. In the case of an EPL attachment circuit the tag is assumed to be a customer tag and it is ignored. In the case of an EVPL attachment circuit the tag is assumed to be a service provider tag and it is popped.

A tagged pseudowire virtual circuit always carries a service delimiting tag. There are two possible options for this type: stamp or push.

- In the case of an EPL attachment circuit the outer tag is assumed to be a C-Tag and an additional tag is pushed on the frame when the C-tag is specified in the vtag-stack attribute when the attachment circuit is defined.
- In the case of an EVPL attachment circuit the outer tag is assumed to be an S-tag and it is stamped, that is, the VLAN tag is altered when the S-tag is specified in the vtag-stack attribute when the attachment circuit is defined.

Figure 2-28 shows how the system determines the pseudowire type.

Figure 2-28
Ingress operation decision tree



Note 1: If identical settings are configured on either end of the MPLS tunnel the original frame is preserved through the service. The operation as the customer frame egresses the service is contrary to the ingress operation. However, the outer tag can also be permanently altered by changing the attachment circuit and virtual switch combination on the egress side.

Note 2: The switch does not behave as indicated in [Figure 2-28](#). L2 transforms need to be configured.

Pseudowire status signaling

Pseudowire status signaling notifies the remote PE about the pseudowire's local forwarding status, which is one of forwarding or not forwarding. A status of not forwarding can be signaled when:

- An operational backup pseudowire is held in fault state until the primary pseudowire is operational. When the primary pseudowire fault is determined, backup pseudowire fault clear is signaled to remote peer.
- Pseudowire is administratively disabled.
- Pseudowire has been manually switched over to the protection pseudowire.

The static pseudowire status signaling relies on explicit acknowledgments from the receiver. The built-in assumption for status signaling over the static pseudowire is that the pseudowire is operational. When it is not, peers need

to exchange the pseudowire status. When node A sends not-forwarding pseudowire status to node B, node B sets a timer to four times the refresh status interval received in the pseudowire status message. If node A does not refresh before the timer expires, then node B reverts the remote pseudowire status to forwarding.

Pseudowire status signaling is sent in band with the user data traffic over the pseudowire OAM channel 0x04 GAL/G-ACh with an ACh type of 0x027.

Virtual circuit connectivity verification profile configuration

Virtual circuit connectivity verification (VCCV) is used to verify the data path connectivity for a pseudowire. VCCV provides a control channel that is associated with a pseudowire, and corresponding operations and management functions, such as connectivity verification, to be used over that control channel. VCCV facilitates transport for fault tools such as LSP ping. VCCV applies to all supported access circuit and transport types currently defined for pseudowires.

You can create a VCCV profile to specify the control channel that is associated with a pseudowire, and to specify corresponding operations and management functions, for example, connectivity verification, to be used over that control channel.

VCCV provides various control channels that follow the same path as pseudowire data. The VCCV packets are intercepted at the egress PE from the control channels and submitted to the control plane for payload processing.

A pseudowire OAM channel is used by VCCV ping/traceroute and static pseudowire status and MAC withdraw signaling. [Table 2-9](#) lists the four CC channel types used by pseudowire OAM, listed according to preference.

Table 2-9
Pseudowire OAM CC channel types

OAM channel	CC channel type
0x01	Control word
0x02	Router alert label (out-of-band)
0x03	TTL exhaust/expiry (out-of-band)
0x04	Associated Channel Header (in-band GAL/G-ACh channel)

Only one channel can be used during the life of the pseudowire. For VCCV profile configuration, the network operator provides the supported channels in the profile and associates this profile for the pseudowire. Error checks are performed during the association with static pseudowire. These include ensuring that none or more than one channel type are selected, and conflicting channel types where pseudowire status signal is enabled but the channel type is other than 4.

The CV payload on the control channel packet is typically ICMP Ping, LSP Ping or BFD packets. If a pseudowire supports control word inclusion in the pseudowire header, the ACH of the VCCV header indicates the type of payload. When a pseudowire is established, the PE negotiates with its peer for the type of control channel and the connection verification tools to use. When a pseudowire is statically configured, the network operator must ensure that VCCV capabilities match at communicating PEs.


Note: If the status TLV is enabled on a pseudowire, VCCV on that pseudowire will always use CC type 4 (GAL/G-ACh) with LSP reply mode even if the pseudowire is created over a static unidirectional TE tunnel. Typically VCCV on a pseudowire over a unidirectional TE tunnel uses CC type 3 with IP reply mode.

Pseudowire reversion

Pseudowire reversion occurs when traffic switches back to the primary pseudowire in a pseudowire protection group after the faults on the primary pseudowire are cleared. The reversion-hold- time delays the switchover after the primary pseudowire is up. The primary pseudowire is the preferred pseudowire. The secondary pseudowire is meant for temporary failures on the primary pseudowire.

Figure 2-29 describes the flow of traffic.

Figure 2-29
Pseudowire reversion



The primary pseudowire transitions to the up state when it becomes available and there are no other pseudowires in the group which are up. Otherwise, the primary pseudowire enters into the reversion pending state until the reversion hold-down timer expires. If a secondary pseudowire goes down while the primary pseudowire is in the reversion pending state, the primary pseudowire becomes active immediately. If any down/up transitions happen on the primary pseudowire while it is in reversion pending, the reversion hold-down time is reset and the waiting period starts again.

The default of a pseudowire protection group is reversion on with the reversion hold-down time set to 30 seconds. A reversion hold-down time of 0 means that there is no delay in switching over to the primary pseudowire. When the reversion is off, the time-based reversion is disabled and switchover operations must be done manually. With reversion set to on, the primary pseudowire is put into the reversion pending state if it is up, but not currently active. Reversion off turns off the reversion immediately even if the primary is in the reversion pending state. The change of the reversion hold-down time does not affect the PWs which are already in the reversion pending state.

If after pseudowire group creation, a standby pseudowire becomes operational before the primary pseudowire, the standby pseudowire is temporarily held in inactive state for less than two seconds to let the primary

pseudowire become operational and active. If the primary pseudowire does not become operational during this period, the standby pseudowire becomes active.


Pseudowire reversion can be overridden and restored by setting the manual switchover to on or off. This can only be done when the target pseudowire is up. The target pseudowire is the secondary pseudowire for on and the primary pseudowire for off. When manual switchover is turned on, the secondary pseudowire becomes the preferred pseudowire resulting in the reversion becoming disengaged and an immediate switchover to the secondary pseudowire. If the secondary pseudowire is the active pseudowire when the manual switchover is turned on, it stays active with the reversion disengaged. In this mode, the pseudowire is still protected. If any faults occur on the pseudowire, a switchover back to the primary pseudowire is performed if the primary pseudowire is up. Regardless of the switchover's success, the manual switchover mode is canceled when the secondary pseudowire goes down. When the manual switchover is turned off, the switchover attempt to the primary pseudowire takes place immediately. The reversion mechanism is re-engaged regardless of the switchover's success.

Multi-segment pseudowire

Multi-segment pseudowire (MS-PW) connects pseudowires between two or more different MPLS networks or control plane domains to form a single, virtually contiguous end-to-end pseudowire. The PE nodes at the two ends of the MS-PW are terminating PEs (T-PEs), while the intermediate PEs are switching PEs (S-PEs). A pseudowire in one MPLS network is switched to a pseudowire in the next MPLS network.

Figure 2-30 shows a sample MS-PW architecture.

Figure 2-30
Sample multi-segment pseudowire architecture



The segment pseudowires form one single pipeline for transporting the user data from one end to another, once they are connected. The pseudowire packet data units are switched from one pseudowire segment to another without changing the pseudowire payload.

Each pseudowire segment can be statically configured or dynamically signaled by means of LDP using FEC 128 or FEC 129. The PWs can, however, only be statically stitched together at each S-PE.

MS-PW is used to:

- span different tunneling technologies, for example, MPLS-TP in the access network and MPLS-TE in the core network. MS-PW can be used to implement MPLS-TE across multiple areas by combining two or more MPLS-TE LSPs.
- span multiple administrative domains, for example, when the access network belongs to one service provider and the core network belongs to another service provider
- scale a large network:
 - access PWs from different locations can be aggregated at the switching point onto a smaller number of protected tunnels between the aggregation network and the core network
 - a large network of VPWS or VPLS services can be scaled
 - a full mesh of T-LDP sessions between VPLS nodes is not required

Switching Point Provider Edge Type, Length and Value (SP-PE TLV) also indicates the location of a fault. For this reason, one single SP-PE TLV can be included in the LDP Notify message for the dynamic pseudowires and in the pseudowire status OAM packet for the static pseudowires along with the pseudowire status which contains faults.

Pseudowire status management

A MS-PW can only pass traffic when all pseudowire segments are up. The MS-PW status is reflected correctly only if “status-tlv on” is configured on every segment of the MS-PW. The pseudowire status is transmitted hop-by-hop along the MS-PW path. When there is a local status change on one of the S-PE segments, the updated pseudowire status OAM frames are sent on both segments to inform the neighbors of the change. The pseudowire status refresh occurs independently on each segment according to its own refresh interval. When the status is not zero, the SP-PE TLV is appended to the pseudowire status frame to indicate the fault location.

- a local transmit fault on one of the segments is translated to a local receive fault on the other segment and vice versa. For example, PW1 and PW2 are stitched together on the S-PE. The local egress tunnel fault out PW1 translates to the local ingress tunnel fault out PW2.

- pseudowire status is transmitted hop-by-hop. This means that at each S-PE, the received pseudowire status on PW1 is merged with the local status of itself and the local status of PW2, and then sends out PW2.
- An admin disabled static pseudowire has the “PW NFwd” fault set in the status to be sent.
- One SP-PE TLV is included in the pseudowire status packet to indicate the location of the nearest fault. There is no SP-PE TLV attached to the pseudowire status frame if no faults are present. The following sub-TLVs in the SP-PE TLV are supported:
 - PW Id: the PW Id of a pseudowire segment where the fault has occurred
 - Local IP: the local IP address associated with the pseudowire segment
 - Remote IP: the remote IP address associated with the pseudowire segment
- Fault to next hop shows VC output on the S-PE to display the overall pseudowire status sent to the next hop. The SP-PE TLV information is also shown if there are faults in the status.

Note: For the static pseudowire status signaling to work properly in the MS-PW environment, the T-PEs have to be upgraded to the releases which support MS-PW.

TTL of the pseudowire label

TTL of the pseudowire label operates as follows:

- Data packets enter the first segment of the MS-PW with the TTL at the pseudowire label set to 255. The TTL is reduced by one at each S-PE.
- For the VCCV profile OAM packets, the TTL is set to the number of pseudowire hops which the OAM packets traverse through and is reduced by one at each S-PE.
- For the static pseudowire status packets, the TTL of the pseudowire label is set to 1 so that the next hop can receive the PW status, and then rebuild and pass it onto the next segment.
- For the static PW Mac Withdraw packets, the TTL of the pseudowire label is set to 255 so that the packets can go through the intermediate S-PEs and reach the remote T-PE.

Fast pseudowire fault reporting

Virtual circuit (VC) protection is supported at the T-PE acting as the multi-tenant unit (MTU). The VC protection is configured as usual at the MTU. The MS-PW configuration is used to configure one MS-PW along the primary path and the other along the secondary path. The same pseudowire protection logic governs the MS-PW switchover.

When a PE detects a local fault, it reports the fault with TTL=1 to its next hop and both of its next hops if S-PE. The PE also reports the fault with TTL=255. This results in the fault reporting OAM packet traveling along the MS-PW path to the T-PEs without being intercepted. The SP-PE TLV is not included in this OAM packet and each PE still relies on the hop-by-hop OAM packet to learn the up-to-date nearest fault location. The pseudowire status OAM packet with TTL=255 is only initiated when a fault occurs. When the fault is cleared, only the TTL=1 OAM packet is generated. Otherwise, if there are other faults along the path, the MTU may perform an unnecessary switchover which causes the pseudowire to flap and traffic loss.

Limitations and compatibilities

Table 2-10 outlines the limitations and compatibilities for MS-PW.

Table 2-10
MS-PW limitations and compatibilities

Limitations and compatibilities	Description
MS-PW type	Static-to-static and static-to-dynamic PWs are supported.
Static pseudowire status	The remote status of a static pseudowire is set to no fault if no status is received for four times the refresh interval. This causes the black holing of the traffic if the remote is actually unprovisioned. For MS-PW, disabling the pseudowire at the remote and the peer pseudowire segment on the S-PE before unprovisioning the pseudowire can help the proper pseudowire status be reported to both T-PEs and prevents the traffic black holing.
Status query of the static pseudowire	To inter-operate with other vendor's devices, the query bit in the pseudowire status frame may have to be turned off all the time.
VCCV types	Since the hardware can only support one VCCV type for each pseudowire and both pseudowire status and MAC withdrawal are desirable for MS-PW, CC type 4 should always be used regardless if the data plane can support CC type 3. Unlike the other pseudowire types, the MS-PWs have to be deleted and re-created to update the VCCV profiles associated with them.

Forwarding Equivalence Class 129

Forwarding Equivalence Class (FEC)-129 is supported as follows:

- FEC-129 is supported on all tunnel types.
- Multi-segment pseudowire (MS-PW) does not support FEC-129 pseudowires.
- FEC-129 is not supported on spoke pseudowires.
- FEC-129 is supported on PW status signaling

Control word for pseudowire

Control word allows routers performing MPLS payload inspection, such as equal-cost multi-path (ECMP) load balancing, to differentiate pseudowire (PW) payload from an IP payload.

For PW to operate correctly, a PW packet cannot appear as an IP packet to a label switching router. When the application using the PW is aware of packets being out of order or where out-of-order packets disturb the handling of the PW, the PW must avoid being subjected to ECMP load balancing.

An IP packet always starts with the first nibble as 0100b or 0110b, which is examined by the router performing MPLS payload inspection. For the PW to work properly and to prevent incorrect processing of the payload carried within the PW, PW payload must not start with the first nibble as 0100b or 0110b.

To assist with this, a generic PW control word header can be used. Its two formats are described in [“Headers” on page 2-46](#).

Headers


When the PW is operating with control word ON, it adds an additional 4-byte control word header just after the PW header. This header helps the LSR to identify whether IP-ECMP can be applied to the PW payload.

PW MPLS control word header

[Figure 2-31](#) shows the format of the data packet after adding the PWMCW header.

Figure 2-31

Packet format with PWMCW header



[Figure 2-32](#) shows the detailed description and header format for PWMCW.

Figure 2-32
PWMCW header format




Table 2-11 describes the PWMCW format.

Table 2-11
PWMCW header format

Header	Description
Bits 0..3	Differ from the first four bits of an IP packet (start with 4 or 6) and provide the necessary MPLS payload differentiation.
Flags (bits 4 to 7)	The current PW implementation does not use these bits and they are set to 0 in the CW header.
FRG (bits 8 and 9)	The current PW implementation does not use these bits and they are set to 0 in the CW header.
Length (bits 10 to 15)	If the MPLS payload is less than 64 bytes, the length field is set to the PW payload plus the length of the PWMCW. Otherwise it is set to 0.
Sequence number (bits 16 to 31)	Implements the sequencing function. The current PW implementation does not use these bits and sets them to 0 in the CW header. If a non-zero sequence number is received, it is ignored.

PW ACH header

The general format of the control packet after adding the PWACH header is shown in [Figure 2-33](#).

Figure 2-33
Packet format with PWACH header




Figure 2-33 shows the detailed description and preferred header format for PWACH.

Figure 2-34
PWACH header format




Table 2-12 describes the PWACH header format.

Table 2-12
PWACH header format

Header	Description
Bits 0..3	Must be 0001b to allow the PW payload to be distinguished from an IP-payload.
Version	Version number of the PWACH. This specification defines version 0.
Reserved	The current implementation does not use these bits. Reserved must be set to 0, and ignored on reception.
Channel type	Defined in the IANA PW Associated Channel Type Registry.

Control word support

Control word is supported for dynamic, static and multi-segment PW. Protection PW does not inherit the control word from the primary PW. Control word must be configured separately if it is required. It is possible that the primary PW can have control word ON, but protection can have it OFF.

If the control word setting has to be changed, the user must admin down the PW at both ends and perform the set operations to achieve the desired behavior.

Table 2-13 describes the control word support on pseudowires.

Table 2-13
Control word support

Pseudowire	Description
T-LDP PW	Ciena supports T-LDP which enables the establishment of a dynamic PW between two routers using targeted LDP signaling. When a dynamic PW is configured to use control word, it sets the control word bit in the FEC information under fec tlv. This fec tlv is carried inside the label mapping message of the T-LDP protocol. When this label mapping message is received from the peer, the control word is negotiated. The final negotiated value is the operating control word for that PW. It is possible that PW can be configured to use control word, but the operating control word is OFF after negotiation. The VCCV capabilities are also advertised inside the interface tlv carried under the same fec tlv of the label mapping message. As part of the negotiation process, the final CC type is identified. If the operating control word state for PW is on, and VCCV have CC-1/3/4 capabilities on local and remote, then CC-1 is selected as the negotiated CC type, but if CC-1 capabilities is not there on local or remote, then the negotiated CC is selected as none. If the operating control word state is off and VCCV has CC-1/3/4 capabilities on local or remote then CC-3 or 4 is selected as the negotiated CC type.
Static PW	There is no negotiation for control word for static PWs. The configured control word state is the operating control word state on that PW. The user is responsible to have the same control word state on both the end points. If the control word is ON, the operating VCCV CC type is 1.
MS PW	If the operating control word state of a segment on SPE does not match with the operating control word state of its other stitched peer, then that segment PW is not allowed to be up and therefore, the whole MS-PW is not UP. When the operating control word is ON and the CC state (negotiated of a segment) on SPE do not match with its other stitched peer, then CC type is marked as none on the segment which does not have CC-1. But MS-PW end2end remains up. To perform set/unset operation on the PW at SPE if the peer segment is present, the user must first detach the PW from the VS and then perform the operation.

VCCV support

A new default VCCV profile is added to support CC type 1. [Table 2-14](#) describes the VCCV profiles supported.

Table 2-14
VCCV profiles

VCCV profile	VCCV capability	Description
Default PWCwVccvProfile	CC-1, CC-3, CC-4	This profile is used when the PW is created with control word ON and uses the default VCCV profile.
DefaultPwVccvProfile	CC-3, CC-4	This profile is used when the PW is created with control word off and uses the default VCCV profile.

User VCCV profile

The cc-in-band parameter is added to the VCCV profile. This parameter supports VCCV CC type 1. [Table 2-15](#) describes the user VCCV profiles supported.

Table 2-15
VCCV profiles

Control word configured	VCCV capability	Support
ON	CC-3, CC-4	This combination is not allowed. An error is generated.
OFF	CC-1, CC-3, CC-4	This combination is not allowed. An error is generated.

MPLS OAM

[Table 2-16](#) describes which VCCV CC types are supported for MPLS OAM.

Table 2-16
MPLS OAM

Control word configured	VCCV capability	Support
ON	CC-1	CC-1 is the only VCCV CC type supported.
OFF	CC-3, CC-4	CC-3, CC-4 are supported.

Table 2-17 describes the various scenarios with control word and CC type for the PW.

Table 2-17
PW control word and VCCV type combinations

Control word state	VCCV capability	Support
On	CC type 1	Support CC type 1
On	CC type 3	MPLS OAM not supported
On	CC type 4	MPLS OAM not supported
Off	CC type 1	MPLS OAM not supported
Off	CC type 3	Support CC type 3
Off	CC type 4	Support CC type 4

Limitations

Integrated routing and bridging (IRB) is not supported with control word.

Control word is supported on the 3916, 3926m, 3928, 3930, 3931, 3932, 5142, and 5160 devices.

IGMP (Internet Group Management Protocol) over MPLS with control word (CW) and more than one MPLS label is not supported on the 3916, 3930, 3931, and 3932 devices.

Troubleshooting

Complementary protocols used for troubleshooting are:

- LSP ping
- LSP traceroute
- VCCV ping
- VCCV traceroute

LSP ping

LSP ping provides the ability to verify connectivity and detect faults of tunnels through the exchange of standard Echo Request and Echo Reply messages.

LSP ping is supported in IP environments and for co-routed tunnels only in non-IP environments.

LSP traceroute

LSP traceroute provides the ability to verify the path and isolate faults of tunnels through the exchange of standard Echo Request (with increment TTL) and Echo Reply (with downstream mapping from transit nodes).

LSP traceroute is supported in IP environments and for co-routed tunnels only in non-IP environments.

VCCV ping

VCCV ping verifies the data path connectivity for a pseudowire through the exchange of ICMP echo frames over an OAM channel. VCC ping has its own OAM channels: 0x25 for non-ip/udp and 0x21 for ip/udp. MAC withdraw and status TLV have different channel types from VCCV ping. For standby, MS-PW ping is supported for static-to-static MS-PW only

For a ping packet to reach its destination, the TTL must be provided. For MS-PW, the segment parameter must be specified. If the segment parameter is not specified, it defaults to 1 or one segment, and the other segment parameters will be automatically filled in by SAOS with the information it knows about the segment to ensure that FEC validation is correct.

FEC validation comprises

- stack validation, which specifies whether the target node of the ping request packet verifies the FEC content in the FEC stack TLV
- reverse stack validation, which enables FEC validation of the ping reply packet. This option is not available for MS-PW.

VCCV traceroute

VCCV traceroute verifies the path and isolates faults in a virtual circuit network. It is supported on single-segment pseudowire (SS-PW) and multi-segment pseudowire (MS-PW). VCCV traceroute can be initiated from any node to any other node within the MS-PW. For standby, MS-PW traceroute is supported for static-to-static MS-PW only.

Each ping that initiates from the requesting node starts with TTL 1 which increments for each subsequent ping until the end node (or segment) is detected. For a SS-PW, the process does not go any further than TTL 1 or one segment. For MS-PW, TTL increments until all the segments have been traced to the specified segment number.

FEC information for the next segment (source IP, destination IP, and pw-id) is returned by each replying node in the MS-PW which the requesting node includes in the ping packet to the next node. When the ping packet is received, each receiving node verifies the information and informs the user if the verification passes or fails.

Remote management for MPLS

You can access an IPv4 Remote Management Interface over an MPLS VPLS mode virtual switch as an alternative to the management VLAN.


Before configuring management over MPLS, you must first allocate resources to the transport-oam feature on 3916/30/31 platforms. For more information, see [“Allocating resources for an MPLS management virtual switch \(3916, 3926m, 3928, 3930 and 3931 platforms\)”](#) on page 15-10.

This capability is supported on the following platforms:

- 3916
- 3926m
- 3928
- 3930
- 3931
- 3932
- 3942
- 5142
- 5150
- 5160

Figure 2-35 shows a sample topology with a remote management over an MPLS VPLS virtual switch.

Figure 2-35
Remote Management Interface over an MPLS VPLS virtual switch



In the remote device at the left of the figure, there is a management virtual switch, that is, a virtual switch that handles the virtual circuits that carry the control messages. One or more MPLS virtual circuits can belong to the management virtual switch. In the remote device, the control messages arrive at the management virtual switch, which is on the switches, and must go to the host CPU for processing. The CPU performs Layer 3 processing as required.

To create a connection between the management virtual switch and the CPU, you create a CPU sub-interface and attach it to the virtual switch, and then you create a remote interface and associate it with the CPU sub-interface.

In order to carry remote interface traffic over an MPLS tunnel, the remote-interface is associated with a virtual switch. One or more MPLS virtual circuits can belong to the management virtual switch, which allows the management traffic to travel over the MPLS tunnel or tunnels.

This implies that management access to the switch can now be gained from any of the members of this virtual switch, including attachment circuit members. As such, if there are customer attachment circuits on the virtual switch that is associated with the remote interface, customers may be able to obtain management access to the node. In order to prevent this, it is recommended that the service provider create an MPLS virtual switch specifically for use for management, and only add a port where the VLAN-based management traffic arrives to the MPLS virtual switch as an EVPL attachment circuit on a boundary node where the network transitions from VLAN-based management to in-band MPLS-based management.

Bidirectional Forwarding Detection

Bidirectional Forwarding Detection (BFD) is a protocol that provides low-overhead, fast detection of link or node failure between any two nodes running over any medium.

- BFD control packets are transmitted in UDP packets with well-known destination port 3784 within an IPv4 packet.
- As an intermediate node of an Multiprotocol Label Switching (MPLS) Label Switch Path (LSP), the 3916, 3926m, 3928, 3930, 3931, 3942, 5142, 5150, and 5160 platforms run BFD over an IPv4 network to the next hop neighbor (RFC 5880).
- As the end node of an MPLS LSP or Label Edge Router (LER), the 3916, 3926m, 3928, 3930, 3931, 3942, 5142, 5150, and 5160 platforms use BFD to support MPLS LSPs. BFD packet content is encapsulated in an MPLS label stack and sent along the same data path as the LSP through multiple hops to its eventual destination (RFC 6428).
- A very fast Continuity Check (CC) detects a loss of connectivity between two end nodes with detection times in the few tens of milliseconds.
- Connectivity Verification (CV) verifies that a 3916, 3926m, 3928, 3930, 3931, 3942, 5142, 5150, or 5160 source node is connected to the desired end node.
- CC and CV failures are reported so appropriate action can be taken (for example, switching to a backup path).

- HW LSP-BFD over an MPLS tunnel running over an L2 LAG port is not supported.
- HW and SW LSP-BFD is mutually exclusive on platforms that support HW LSP-BFD.
- HW BFD is supported on the 3928, 3942, 5142 and 5160.
- BFD needs to be disabled globally before changing modes.
- Connection failure is not supported.
- IP BFD is not supported in hardware.

MPLS-BFD implements a single-session BFD over an associated bidirectional LSP, constructed from a pair of unidirectional LSPs, one in each direction. The two LSPs operate virtually as one single entity for the purpose of protection and switching.

The system supports the creation of SW VCCV BFD sessions using CC-Type 1 or CC-Type 4 PWs that meet the following requirements:

- Each of the PE nodes for the PW must be one of the following platforms: 3916, 3930, 3931, 3932.
- The PW must be configured statically over an IPv4 network. The PW can be an MPLS Ethernet single segment (SS-PW) configured statically, or an MPLS Ethernet multi-segment (MS-PW) in which each segment is configured statically.
- The PW must be configured over a static co-routed tunnel.
- The PE nodes at the ends of the PW must be configured in the Active role.

There is a maximum of one VCCV BFD session per PW. The BFD session for a qualifying PW is created only when the PW is attached to a virtual switch, and only if BFD monitoring has been enabled on the PW.

For examples showing how to configure VCCV BFD on a PW, see [VCCV BFD over static SS-PW over single LSP example](#) and [VCCV BFD over static MS-PW over single LSP example](#).

A BFD session begins with the periodic, slow transmission (one second intervals) of control packets between nodes. In the initial message exchanges, the nodes exchange state information and their respective discriminator fields (which identifies the session) and negotiate the desired transmit and receive intervals. Subsequent control packets sent from each side must reflect the discriminator values back to the originating node in the Your Discriminator field.

Once two-way communication is achieved and the session is in the Up state, periodic BFD control packets are transmitted at the transmit interval negotiated in the initial packet exchange, and each packet either updates or reiterates the session state information to the other node.

BFD sessions

BFD sessions monitor connections. A BFD session for a static co-routed MPLS-TP tunnel monitors the connection between the two ends of the tunnel. BFD sessions can also be configured to monitor the connection between adjacent network elements.

Start-up of a BFD session

When a BFD session begins, the transmit interval is not yet determined, so the periodic transmission of control packets between nodes occurs in one second intervals. In the initial message exchanges, the nodes exchange state information and their respective discriminator fields (which identifies the session) and negotiate the desired transmit and receive intervals. Subsequent control packets sent from each side must reflect the discriminator values back to the originating node in the Your Discriminator field.

Once two-way communication is achieved and the session is in the Up state, periodic BFD control packets are transmitted at the transmit interval negotiated in the initial packet exchange, and each packet either updates or Defect handling.

Once BFD detects a failure, the failure is reported to the relevant protocol module which can then initiate the appropriate action, as specified by their respective protocols. The tunnel is not disabled unless the next hop is unreachable. If the next hop is reachable, the tunnel remains operationally and administratively up and its tunnel role is changed to standby.

Only the timeout and path down failure conditions cause a switch to the backup path. Control packets received with other diagnostic codes (such as Administratively Down) may transition the BFD state to “Down” state but will not cause a switch to the backup path.

Defect criteria and handling

Three things will cause a transition to the defect state:

- BFD session times out after not having received BFD CCs 3.5 times in a row
- Receipt of a Interface Path Down notification
- Connectivity Verification (CV) failure

CV fails when the following checks fail:

- Invalid encapsulation, such as receiving IP/UPD encapsulated CC or CV packets on an LSP that uses the Generic Associated Channel (G-ACh), as described in RFC 6428.
- Receipt of an invalid source MEP-ID TLV
- Receipt by the local node of a BFD CC or CV packet with a non-existent Your Discriminator value while the session is in the Up state
- Receipt of CC or CV packets with an invalid label

When a node detects a failure, each node initiates a switch to the backup LSPs if they exist.

Maximum number of BFD sessions

The BFD scale numbers mentioned in this section are one-dimensional. They do not consider other configured features or BFD flavors.

Table 2-18 lists the maximum number of IP BFD sessions by platform.

Table 2-18
Maximum IP BFD sessions

Interval	IP BFD Maximum		
	3942/5142/5160	3926m/3928/3930/ 3932/5150	3916/3931
3.3 ms	--	--	--
10 ms	32	24	16
100 ms	100	100	50
300 ms	200	200	100
1 sec	less than or equal to 200	less than or equal to 200	less than or equal to 100
10 sec	less than or equal to 200	less than or equal to 200	less than or equal to 100

Table 2-19 lists the maximum number of LSP BFD sessions by platform.

Table 2-19
Maximum LSP BFD sessions

Interval	LSP BFD Maximum									
	3942/5142/5160		3930/3932		5150		3916/3931		3926m/3928	
	SW	HW	SW	HW	SW	HW	SW	HW	SW	HW
3.3 ms	N/A	380	N/A	N/A	N/A	N/A	N/A	N/A	N/A	200
10 ms	32	750	24	N/A	24	N/A	16	N/A	24	200
100 ms	100	1000	100	N/A	100	N/A	100	N/A	100	200
300 ms	200	1000	200	N/A	200	N/A	100	N/A	200	200
1 sec	1000	1000	300	N/A	400	N/A	200	N/A	300	200
10 sec	1000	1000	300	N/A	500	N/A	300	N/A	300	200

Table 2-20 lists the maximum number of VCCV BFD sessions by platform.

Table 2-20
Maximum VCCV BFD sessions

Interval	VCCV BFD Maximum									
	3942/5142/5160		3930/3932		5150		3916/3931/3960		3926m/3928	
	SW	HW	SW	HW	SW	HW	SW	HW	SW	HW
3.3 ms	N/A	310	N/A	N/A	N/A	N/A	N/A	N/A	N/A	270
10 ms	32	600	24	N/A	N/A	N/A	16	N/A	24	800
20 ms	48	735	36	N/A	N/A	N/A	24	N/A	36	128
50 ms	64	765	64	N/A	N/A	N/A	50	N/A	64	128
100 ms	100	800	100	N/A	N/A	N/A	100	N/A	90	915
300 ms	200	800	200	N/A	N/A	N/A	100	N/A	200	915
1 sec	800	800	300	N/A	N/A	N/A	200	N/A	300	915
10 sec	800	800	300	N/A	N/A	N/A	300	N/A	300	915

BFD mode

BFD can run in two modes: asynchronous or demand mode. The system software supports asynchronous mode. When BFD runs in asynchronous mode, Continuity Check sends and receives periodic control packets that verify the integrity of the link or nodes. Link or remote node failures are

detected once control packets have been lost for three consecutive receive intervals. BFD mode times out after not having received BFD CCs 3 1/2 times in a row.

BFD roles

Either node can be configured to take an Active or Passive role. At least one of the nodes must be configured as Active. The node taking the Active role initiates the session establishment procedure, regardless of whether or not it has received any BFD packets for the session. The node configured to be Passive will not begin sending any control packets until it has received a control packet. Once the BFD session is established, both nodes will send and respond to BFD control packets regardless of their role.


Note: One or both of the nodes must take the Active role in order for the BFD session to be established. If both nodes are configured in the Passive role, the BFD session will not be established.

BFD control packet format

Figure 2-36 shows both the optional and mandatory portions of a BFD packet format.

Note: Currently, BFD does not support authentication.

Figure 2-36
BFD control packet format



BFD state machine

Figure 2-37 shows the BFD state machine and the three states through which a BFD session normally proceeds:

- Init
- Up
- Down


A fourth state, AdminDown, exists to administratively hold down the session indefinitely thereby causing the remote system to enter the Down state.

Each system communicates its session state information in the State field of the control packet.

- 1 A typical BFD session establishment scenario starts off with both systems in the Down state.
 In this state, the system that takes the Active role will send the remote system a control packet with this state information which will cause the remote system to advance to the Init state.
- 2 The next control packet received from the remote system (with the state set to Init) validates the two-way communication and brings the session of the local system to the Up state.
- 3 The session state of the remote system will also advance to the Up state on receipt of the next control packet from the local system.

- 4 In the Up state both systems send periodic control packets to each other if operating in the Asynchronous mode.
- 5 Any time control packets have been lost for Detect Mult times in a row, the session is taken down. Detect Mult is the detect-multiplier value, which is the number of packets that must be missed consecutively to declare the session to be down. The value is sent by the far-end system.

Figure 2-37
BFD state machine




Single-hop IPv4 BFD

Figure 2-38 shows how BFD can be configured to run over an IPv4 interface but only to its immediate neighbor node in the same subnet. This single-hop, IP-based configuration can be used to detect connectivity failures between a pair of adjacent LSRs, which make up one section of a multi-section LSP.

IS-IS linkage to BFD allows an IS-IS adjacency to be tied to the IP BFD sessions on an interface, which accelerates error detection and allows for the traffic engineering database to converge quicker and ultimately improve the dynamic tunnel convergence time.

Figure 2-38
Single-hop IPv4 BFD



BFD sessions

Only one BFD session is needed between the two nodes. One node must be configured to run in the Active role while the other must be configured as Active or Passive. Once the BFD session is successfully established, that is, initially established with default transmit and receive intervals of one second, both intervals are changed to the values configured on the CLI.

Defect handling

For IP-BFD, the node requires a next-hop IP neighbor with a resolved ARP. Fault notification withdraws this IP adjacency, which is monitored by the MPLS LSPs that use it as a next-hop, and they fault as a result.


For LSP BFD, an LSP BFD session can be established over a static MPLS-LSP if the ingress LSP has been associated with an egress LSP.

Once BFD detects a failure, the failure is reported to the relevant protocol module which can then initiate the appropriate action, as specified by their respective protocols.

Associated bidirectional LSPs

Figure 2-39 shows how associated BFD in asynchronous mode runs over a pair of unidirectional LSPs, that is, one in each direction. These unidirectional LSPs are associated or paired with one another at the ingress/egress points. One single BFD session is used for the LSP pair, operating in coordinated mode where the session state tracks the defect state.

Figure 2-39
Associated bidirectional LSPs



The two LSPs are bound by means of static configuration using an MPLS CLI command at both the tail- and head-ends of the LSP. The selected pair must be the same pair at each end, otherwise the BFD session cannot be established. The originating LER is configured in the Active role and the remote LER is configured in the Passive role. You can also configure the originating LER in the Passive role and the remote LER in the Active role.

Note: Although configuring both nodes as Active will also work, configuring both as Passive is a misconfiguration and the BFD session will not be established.

Periodic CC and CV messages are sent in both directions, that is, ingress to egress and egress to ingress, once the session is up and running. The CCs are sent at the negotiated transmit interval while the CVs are sent once every second.

Defect criteria and handling

Three things will cause a transition to the defect state:

- BFD session times out after not having received BFD CCs three times in a row
- receipt of a Interface Path Down notification
- Connectivity Verification (CV) failure

CV fails when the following checks fail:

- Invalid encapsulation, such as receiving IP/UPD encapsulated CC or CV packets on an LSP that uses the Generic Associated Channel (G-ACh), as described in RFC 6428.
- Receipt of an invalid source MEP-ID TLV
- Receipt by the local node of a BFD CC or CV packet with a non-existent Your Discriminator value while the session is in the Up state
- Receipt of CC or CV packets with an invalid label.

Faults detected by missing CVs follow the same frames being missed for 3.5 times in row, that is, 3.5 seconds for CV.

Interoperability

Ciena's BFD supports G-ACh type 0x0022 for CC BFD as described in RFC 6428. In addition, Ciena's implementation supports an additional option in the BFD profile for compliance with existing BFD implementations. When the BFD profile is created you can set the `lsp-gachtype` attribute to 7 or 22:

- set to 7, BFD will do CC
- set to 22, BFD will do CC and CV

VCCV BFD

VCCV BFD can be configured to run over an MPLS PW and monitor the PW service for receive and transmit faults.

Notes

- VCCV BFD is supported on MPLS Ethernet Single- and Multi-Segment (with individual PW segments configured statically) PWs only, configured over static co-routed tunnels only, over an IPv4 network.
- VCCV BFD is supported only on PW end-nodes (TPEs/MTU). A maximum one VCCV BFD session per PW is supported.
- BFD is only supported on PWs configured with VCCV CC-Type 1 (PW ACH) and CV-Type 0x10 (without IP/UDP headers and for Fault Detection only).
- A BFD session is created only when BFD-monitoring is explicitly enabled on a PW and the PW is attached to a VS.
- An existing BFD session is deleted when either BFD-monitoring is disabled on the PW or the PW is detached from the VS.
- A VCCV BFD session is moved to an Admin-Down state before session deletion to prevent false fault notifications on the remote session.
- BFD fault detection starts only after the VCCV BFD session becomes operationally up (that is, the BFD state becomes Up). VCCV BFD does not detect network faults in Down state.
- A BFD session changes to an administratively down state if the associated PW is disabled, that is, changes to Admin-Down (RFC 5885). Fault detection is stopped when local or remote PW is disabled.
- Any existing VCCV BFD faults are cleared in the following cases:
 - When BFD-monitoring is disabled or VS is detached at the local or remote ends.
 - When PW is disabled at the local or remote ends.
 - When the BFD session becomes operationally up again.

- The creation of a VCCV BFD session on PW is independent and agnostic of creation of LSP BFD sessions on underlying tunnel. The user can configure proper BFD timers (as per actual service/network latency) to achieve “Fault Hierarchy”, thus propagating the BFD fault from LSP till PW, such that PW switchover starts only after either LSP switchover gets completed or MBB gets completed successfully.
- VCCV and LSP BFD session resources are shared on both hardware and software. The maximum number of VCCV and LSP BFD sessions combined remains unchanged.

Limitations

- VCCV CC-Type 4 is not supported on 3926m, 3928, 3942, 5142, and 5160 platforms due to BCM XGS chipsets not supporting PW-GAL for VCCV BFD.
- VCCV BFD sessions can be created on software or, if the product supports it, on hardware. Simultaneous co-existing VCCV BFD sessions on both hardware and software are not supported. However, VCCV BFD on hardware can co-exist with LSP BFD on software and vice-versa.

BFD dampening

In networks, dynamic network events (such as link flapping) can lead to BFD session operational state toggling. In normal circumstances, every BFD session operational state change is notified to the BFD client application and SNMP server. If BFD session flaps (that is, operational state changes from Up to Down) occur too frequently, BFD notifications can overload the client applications. This may degrade the overall system performance.

Usually there is no restriction on BFD sessions to notify session state changes to clients or to generate traps to the SNMP server. If the flaps occur too frequently and continue for a while, they may cause the SAOS system and/or SNMP server to overload, causing the system to become unresponsive and unavailable. To mitigate such system overloading behavior, Ciena recommends to suppress the client notifications and traps simultaneously on a per-session basis when the session flaps are encountered.

The suppression begins when the session encounters a given number of flaps and finally goes Down. When the session becomes stable again (that is, the session continues to be operationally down or returns to an operationally up state) for some time, client notifications and traps are simultaneously re-enabled for the session.

SAOS BFD implements an “Exponential Decay” algorithm to provide a flexible framework for disabling BFD notifications/traps when the network behavior is deemed as too flappy and re-enabling them when the network becomes more stable. This algorithm is similar in behavior to standard BGP route flap dampening guidelines mentioned in RFC 2439 (BGP Route Flap Damping).

The algorithm's behavior is tightly governed by a set of parameters, mostly user-configurable, to allow the user to control the resultant dampening effects as per the dynamic characteristics of the network, such as configuration, topology, network stability, and so on.

Exponential Decay algorithm

The "Exponential Decay" algorithm works on calculations done on a "penalty" metric (maintained per BFD session). All its parameter values are processed as integer calculations. This table captures all of the algorithm parameters and their details.

Table 2-21
Exponential Decay algorithm parameters

Parameter name	Parameter meaning	User configurable	Dynamic range possible	Default value	Calculation logic
penalty	Used to derive the stability of the BFD session.	No	0 – max-penalty	0	Incremented by "penalty-increase", up to "max-penalty" value, every time a flap is detected. Decrement by half after the expiry of periodic "decay-half-life" timer.
penalty-increase	Value by which "penalty" is incremented (when a session flap is detected).	No	1 – (suppress-threshold + 1)	100	None
max-penalty	Ceiling value of "penalty".	No	(As per calculations)	(As per calculations)	$\text{reuse-threshold} * 2^{\frac{\text{max-suppress-time}}{\text{decay-half-life}}}$
suppress-threshold	Threshold value above which notifications/traps are suppressed.	Yes	$\text{reuse-threshold} - (\text{max-penalty} - 1)$	400	None

Table 2-21
Exponential Decay algorithm parameters

Parameter name	Parameter meaning	User configurable	Dynamic range possible	Default value	Calculation logic
reuse-threshold	Threshold value below which notifications/traps are re-enabled.	Yes	1 – suppress-threshold	15	None
decay-half-life	Waiting time after which “penalty” is halved, provided the session state remains stable.	Yes	500 msec - 1,800,000 msec	1,000 msec	Max value denotes approx. 30 min timer assuming minimum 3 sec fault detection & recovery of a BFD session.
max-suppress-time	Maximum time period for which notifications/traps can be suppressed. After this period, notifications/traps are re-enabled, no matter the session stability or state.	Yes	(decay-half-life + 1) msec - 3,600,000 msec	5,000 msec	Max value denotes approx. 60 min timer assuming minimum 3 sec fault detection & recovery of a BFD session.

Algorithm working example

Consider an LSP BFD session (with Rx/Tx intervals configured as 3.3 ms) with these parameter values configured by the user:

- suppress-threshold = 110
- reuse-threshold = 50
- decay-half-life = 2 sec
- max-suppress-time = 11 sec


The internal parameters are:

- penalty-increase = 30 (value assumed for representation purposes)
- max-penalty = $50 * 2^{(11 / 2)} = 2272$

Figure 2-40 provides a graphical view of the different parameters, depicting changes in their values, according to the ongoing network behavior over time.

Note: In this graph, the time units are taken in seconds and latency arising due to BFD protocol processing itself has been ignored.

Figure 2-40
Graphical view of algorithm working example



Notes

BFD dampening supports the following:

- Suppressing/re-enabling BFD notifications to internal SAOS client modules.
- Suppressing/re-enabling SNMP trap/syslog “BFD Oper State Change”.
- A new SNMP trap “BFD State Flap Dampening” for suppression of BFD notifications. The trap indicates that suppression has been enabled (due to session flapping excessively) or disabled (due to session being stable) for a session. The parameters for this trap are:
 - Session name
 - Session type—LSP BFD.
 - Service name—Tunnel name for LSP BFD.

- Session oper-state—Shows current BFD operational state of the session.
- Session admin-state—Shows current BFD administrative state of the session.
- Suppression state—(enabled/disabled)—Shows if client notifications and BFD operating state change traps/syslogs are currently enabled or suppressed.
- Reason—Reason due to which this trap is generated (global-admin-down / session-admin-down/session-delete / global-dampening-disabled / session-dampening-disabled/suppress-threshold / reuse-threshold / max-suppress-timeout/user-trigger / debug).
- Penalty—Shows current penalty metric value.
- Date and Timestamp—Shows date and time when suppression is enabled or disabled.
- A new SNMP syslog “BFD State Flap Dampening” for suppression of BFD notifications. The syslog indicates that suppression has been enabled (due to session flapping excessively) or disabled (due to session being stable) for a session. The parameters for this syslog are:
 - Session name
 - Session oper-state—Shows current BFD operational state of the session.
 - Session admin-state—Shows current BFD administrative state of the session.
 - Suppression state—(enabled/disabled)—Shows if client notifications and BFD operating state change traps/syslogs are currently enabled or suppressed.
 - Reason—Reason due to which this trap is generated (global-admin-down / session-admin-down/session-delete / global-dampening-disabled / session-dampening-disabled/suppress-threshold / reuse-threshold / max-suppress-timeout/user-trigger / debug).
 - Penalty—Shows current penalty metric value.
 - Date and Timestamp—Shows date and time when suppression is enabled or disabled.
- This “BFD State Flap Dampening” trap follows user configuration to enable/disable the trap to be generated to a SNMP server.
- An LSP BFD session for static bidirectional co-routed LSPs.
- All the LSP BFD sessions that can be configured.
- Use of a “penalty” metric (having an integer value) as a measurement criterion for predicting stability of a BFD session.

- These parameters to guide BFD dampening mechanism for each session:
 - default-penalty—Default penalty value which is the minimum value attainable. The value is fixed as 0.
 - penalty-increment—Constant value by which penalty is incremented every time a session flaps. The value is fixed as 100.
 - suppress-threshold—User-configurable threshold value which when crossed above shall enable suppression.
 - reuse-threshold—User-configurable threshold value which when crossed below shall disable suppression.
 - decay-half-life—User-configurable time duration after which the penalty metric is halved. This timer is reset after each flap.
 - max-suppress-time—User-configurable maximum time duration (starting from the time suppression is enabled) for which a session can be suppressed. All notifications are re-enabled after expiry of this time period, no matter the session's prevailing stability.

Note 1: The relationships between the parameters are as follows: (1) default penalty < reuse-threshold <= suppress-threshold < max-penalty; (2) decay-half-life < max-suppress-time.

Note 2: These parameters are consistent with the suggested parameters in RFC 2439 (BGP Route Flap Dampening). However, unlike RFC 2439, the same values of decay-half-life are used for both operationally up and operationally down sessions to reduce the configuration and algorithm complexity.

- Increment the penalty metric each time a BFD session flaps, until it reaches the “max-penalty” ceiling value. The max-penalty is non-configurable and calculated internally as:

$$\text{Max-penalty} = \text{reuse-threshold} * 2^{(\text{max-suppress-time}/\text{decay-half-life})}$$
- Decrement the penalty metric after every “half-life” time interval until it reaches a floor value (default fixed penalty value), provided the BFD session remains stable (that is, the session operating state remains unchanged) during this time interval. This rate of decrement of penalty follows an Exponential Decay rate.
- For a session when the associated penalty metric becomes greater than the “suppress-threshold” value, support suppressing BFD client notification and SNMP trap/syslog for BFD operating state change.
- For a suppressed session when the associated penalty metric becomes lesser than the “reuse-threshold” value, support re-enabling BFD client notification and SNMP trap/syslog for BFD operating state change.
- Creating, modifying, and deleting a BFD dampening profile through CLI, containing these parameters:

- suppress-threshold
- reuse-threshold
- decay-half-life
- max-suppress-time
- mode (test | full)—“Test” mode only calculates and stores the results for a given set of configured dampening parameters over a time period, without doing any real suppression of the sessions. The trap/syslog “BFD State Flap Dampening” is not generated. “Full” mode additionally supports suppressing and re-enabling sessions.

Note: The user must make sure that the dampening parameters configured manually are identical on both BFD endpoints, to ensure synchronous behavior. As mentioned in RFC 7196 (Making Route Flap Dampening Usable), “test” mode is intended to be used only for getting the results of BFD dampening behavior in a topology, according to user-configured dampening parameters, without impacting normal BFD notifications.

- A default BFD dampening profile, “Dampen-Default”, that cannot be modified:
 - suppress-threshold = 400
 - reuse-threshold = 15
 - decay-half-life = 1,000 msec
 - max-suppress-time = 5,000 msec
 - mode = test
- The same parameter values (from “Dampen-Default”) are also used as default values when some of the parameters are not provided.
- The default BFD dampening profile is used when no explicit user-defined profile is provided.
- A maximum of 33 BFD dampening profiles, including one default profile and 32 custom profiles.
- Associating a dampening profile with a new or existing BFD session:
 - BFD monitoring on a service must be enabled before enabling dampening of that session. Consequently, BFD dampening on a session is automatically disabled if BFD monitoring is disabled for the service.
 - BFD monitoring on a service need not be disabled for dampening to be enabled or disabled for that session.

- A custom BFD dampening profile can be modified only if dampening is disabled, and only if the BFD dampening profile is not associated with any BFD session.
- Deleting a dampening profile only if it is not associated with any BFD session.
- The user configuring through CLI to enable or disable the dampening feature, on a global basis (that is, for all BFD sessions) and on a per-session basis.
- The user invoking through CLI to retrieve dampening statistics for a BFD session.
- The user invoking through SNMP/MIB to retrieve dampening statistics for a BFD session.

Alarm Indication Signal with Link Down Indication

Alarm Indication Signal with Link Down Indication (AIS/LDI) messages are in-band Operational, Administration and Maintenance (OAM) messages sent to indicate service-interrupting faults on MPLS LSPs.

AIS/LDI can help minimize service disruption and facilitate problem identification and resolution in a number of ways:

- AIS/LDI helps identify the node on which a fault is detected. Information available at the downstream Maintenance End Point (MEP) identifies the tunnel on which the AIS/LDI message was received along with node and interface on which the fault was detected.
- AIS/LDI can trigger protection switching.

Operation

AIS/LDI fault OAM messages are generated when faults on layers below the LSP are detected by an intermediate node where an LSP is switched.

AIS/LDI monitors previous hop adjacencies for faults using IP-BFD or ARP timeout.

- If IP-BFD is enabled in the transit router, then a failure will result in the IP-BFD session going down. This will cause the IP adjacency to be removed which will then trigger a fault.
- If IP-BFD is not configured on the transit router, then ARP timeout is used to perform the IP adjacency check and indicate that the previous hop has failed. Again, this will bring the adjacency down and trigger a fault.

When a link or port used by an LSP is faulted, an AIS/LDI message is sent towards the end point of the LSP. The message is sent downstream on the LSP that was being used to send data by the node closest to the upstream fault.

Note 1: The system does not wait to send an AIS/LDI message. The message is sent immediately after a hold timeout of zero.

Note 2: AIS/LDI does not distinguish between an IP interface taken down administratively and an IP adjacency taken down because of a fault. Taking an IP interface down administratively generates an AIS/LDI message.


The receiver of an AIS/LDI message uses the signal to trigger protection switching.

When an AIS/LDI message is not received by the downstream endpoint for a period corresponding to 3.5 times the configured refresh timeout value, a fault is assumed to be cleared.

Remote Defect Indication

Figure 2-41 shows how a failure triggers an AIS/LDI message, which in turn triggers a Remote Defect Indication (RDI) message.

Figure 2-41
AIS/LDI and RDI messages



AIS-RDI

If BFD is not enabled on a tunnel, an LER sends a Ciena proprietary message called AIS-RDI. This message is sent separately from any BFD-RDI message in support of systems without BFD. AIS-RDI is enabled by default.

AIS-RDI removes the need for BFD to carry RDI. Instead the AIS/LDI message received by the endpoint is used to generate another message called AIS-RDI and used to send the RDI notification to the upstream MEP.

An AIS-RDI message uses the MPLS-TP fault message format.

- The Ach channel type is 0x58, indicating fault management.
- The message type 0x01, indicating Alarm Indication Signal.
- The Remote Defect Indication Flag is set.
- IntfNum in the Interface TLV is zero.

Because the frame appears as a normal MPLS frame on a particular LSP, it is switched by all intermediate routers. If the upstream MEP is a Ciena product running AIS-RDI, it recognizes the frame as an AIS-RDI frame (because IntfNum in the Interface TLV is zero) and takes the appropriate action. All other vendor MEPs drop the AIS-RDI frame because IntfNum in the Interface TLV is zero.

BFD RDI

If BFD is enabled on a tunnel, an RDI message is sent by an LER on the BFD session immediately after receiving an AIS/LDI message from an upstream router. The MEP determines the associated tunnel of the failing tunnel and uses a BFD session running on that tunnel to send an RDI.


Instead of waiting for the next scheduled BFD packet, an RDI-enabled BFD packet is sent immediately and repeatedly until the downstream MEP receives an AIS/LDI message indicating that the fault is cleared.

Co-routed tunnels

Figure 2-42 shows a link failure in the middle of a co-routed tunnel configuration detected on both unidirectional tunnels. Each failure-detecting router informs its downstream MEP about the failure. This can trigger protection switching. If an RDI message is generated by the downstream MEP, it is dropped in the middle of the network where the fault occurred.

Note: If the downstream MEP determines that intermediate routers are using a co-routed tunnel then it will not send an RDI message.

Figure 2-42
Co-routed tunnels




Associated tunnels

Figure 2-43 shows a failure in the middle of an associated tunnel configuration which causes an AIS/LDI message to be generated by the closest downstream router. This AIS/LDI message flows towards the downstream MEP where it can be used to trigger protection switching.

The downstream MEP then triggers an RDI message to be sent on the associated tunnel towards the upstream MEP.

Figure 2-43
Associated tunnels




Multiple fault conditions

AIS/LDI can deal with multiple fault conditions present in a network simultaneously.

Figure 2-44 illustrates a single fault (F1) detected between P1 and P2.

Figure 2-44
Single fault (F1) detected on the network




As with co-routed tunnels, P1 and P2 send LDI messages faulting the LSP to PE1 and PE2, respectively:

- PE1 shows the fault location as P1.
- PE2 shows the fault location as P2.

Figure 2-45 illustrates a second fault (F2) detected subsequently between P2 and P3.

Figure 2-45
Second fault (F2) detected on the network




As soon as PE2 receives the LDI message sent by P3 to indicate fault F2, PE2 updates the fault location from P2 to P3.

The fault location shown by PE1 is not updated since the LDI messages sent by P2 to inform PE1 about the new fault F2 are blocked by fault F1.

If fault F1 is cleared first, P1 sends a fault clear message to PE1, clearing the fault on PE1. P2 sends a fault clear message to PE2, but the message is blocked by fault F2, and PE2 remains in a faulted condition.

As shown in [Figure 2-46](#), P2 continuously sends LDI messages to PE1 because of fault F2. As soon as PE1 receives this LDI message, it faults the LSP again reflecting the new fault location as P2.

Figure 2-46
First fault (F1) cleared on the network



Note: There will be a window of time during which PE1 will not have a fault on an LSP; however, PE2 will have a fault on that LSP.


Node protection

AIS/LDI should not be used in node protection configurations similar to the one illustrated in [Figure 2-47](#).

When a fault is detected, the first P router finds an alternate path through the bottom P router to re-route traffic. However, the top middle P router detects a fault and immediately sends an AIS/LDI message towards the downstream

LER. This causes the downstream LER to assume that the path including the protected node has actually failed and can trigger protection switching when protection switching is not required.

Figure 2-47
Node protection



AIS/LDI and client/server layers

The link interface layer and hierarchical tunnels interact with client/server layers.


This section describes:

- “Link interface layer faults” on page 2-78
- “Hierarchical tunnel faults” on page 2-79

Link interface layer faults

When a fault is detected at the Link/Interface layer as shown in Figure 2-48, an AIS/LDI message is generated by the network element that detects that fault. Depending on what other fault detection mechanisms are in place, alarms are generated by layers above. For example, BFD would detect and raise an alarm at the LSP level.

Figure 2-48
Link/Interface layer and client/server layers



Since the BFD fault is detected at the downstream MEP only, alarms are normally raised at the downstream MEP.


The upstream MEP does not detect a fault and therefore does not generate an alarm.

Hierarchical tunnel faults

In the case of hierarchical tunnels, the inner LSP is a server for AIS/LDI and the outer LSP is a client of AIS/LDI as shown in [Figure 2-49](#).

Figure 2-49

Hierarchical tunnels and client/server layers



When a fault occurs on the inner LSP, two SNMP traps are sent and protection switchover is initiated on the LER if a protected tunnel is available.

Interworking with other vendor products

AIS/LDI interoperates with other vendors by supporting the GAL/GACh channel for communication of AIS/LDI. [Table 2-22](#) describes AIS/LDI interworking with other vendor products.

Table 2-22
AIS/LDI interworking

AIS/LDI action	Ciena LER	Ciena LSR	Other vendor LER	Other vendor LSR
AIS/LDI	Will receive. Will send RDI if fault is detected.	Will send.	Expected to receive.	Expected to send.
BFD RDI	Will send/receive.	Will switch transparently.	Expected to receive.	Expected to switch transparently.
AIS/LDI RDI	Will send/receive.	Will switch transparently.	Not expected to send. Not expected on receive.	Expected to switch transparently.

Features not available on 393x and 3916 platforms

The following features are not available on 393x and 3916 platforms:

- Border Gateway Protocol (BGP) features (seamless MPLS etc)
- IP over MPLS device management
- IP Bidirectional Forwarding Detection (BFD) – multi-hop
- Fast Reroute (FRR)
- MPLS OAM penultimate hop popping (PHP)
- Make before break (MBB)
- Diffserv-TE
- Label switched path (LSP) re-optimization
- Performance monitoring (PM) statistics
- Auto-bandwidth adjustment and trap
- Pseudowire (PW) control word
- IGMP (Internet Group Management Protocol) over MPLS with penultimate hop popping (PHP), with or without control word (CW)

Initial configuration

In order to configure MPLS features, you need to install the MPLS license key for each line module that MPLS will be used on. License keys can be purchased by contacting Ciena Customer Support.

MPLS implementation requires an L3 IP interface for handling MPLS control protocol traffic. This IP interface is associated with an L2 flood domain.

For more information about logical interfaces, refer to *39XX/51XX SAOS 6.16 Advanced Ethernet Configuration*.

This chapter provides the following initial configuration procedures:

- [“Installing the MPLS license on 39XX/51XX” on page 3-2](#)
- [“Configuring IP for MPLS interfaces on 39XX/51XX” on page 3-4](#)

Procedure 3-1

Installing the MPLS license on 39XX/51XX

You can install a premium feature license key directly by identifying the license key and module number. When the module number is left unspecified, the value defaults to 1.

Step	Action

1	Install a premium feature license key: <pre>software license install [file <String>] [server <IP Address or host name>] [license-key <String[14]>] [module <NUMBER: 1..3>] {default-server default-ftp-server default-tftp- server default-sftp-server {tftp-server <ip-host-str> [server-port <INTEGER: 1...65535>]] {ftp-server <ip-host-str> [login-id <username> [<password-attr> <encrypted-password-attr> <echoless- password-attr>][server-port <INTEGER: 1...65535>]] {sftp-server <ip-host-str> login-id <username> {<password-attr> <encrypted-password-attr> <echoless- password-attr>][server-port <INTEGER: 1...65535>]]}}</pre>

where

file <String>	is the license filename and path.
server <IP Address or host name>	is the TFTP server.
license-key <String[14]>	is the license key string.
module <NUMBER: 1..3>	is the module number.
default-server	uses the default xFTP server.
default-ftp-server	uses the default FTP server.
default-tftp- server	uses the default TFTP server.
default sftp- server	uses the default SFTP server.
tftp-server <ip- host-str>	is the tftp-server.

where

server-port	is the server-port number.
<INTEGER: 1...65535>	
ftp-server <ip- host-str>	is the ftp-server name.
login-id	is the FTP/SFP username.
<username>	
password-attr	enters the password in clear text.
encrypted-pass- word-attr	sets the password using a pre-encrypted string.
echoless- password attr	engages an echoless password collector.
sftp-server <ip- host-str>	is the sftp-server name.

—end—

Example

The following example installs a license key with implied module 1:

```
software license install license-key W123XYZ123XYZY
```

The following example installs a license key with module 2:

```
software license install license-key W123XYZ123XYZY module 2
```

Procedure 3-2

Configuring IP for MPLS interfaces on 39XX/51XX

For each MPLS deployment scenario, an IP interface and loopback interface must be configured for the L2 routing domain to handle MPLS control protocol traffic.

Note: IP interfaces are required to configure MPLS tunnels. This procedure must be completed prior to configuring MPLS tunnels.

IP/MPLS and MPLS-TP tunnels can share the same MPLS interface.

If the switch configuration includes Fast Reroute (FRR), enable IMPLICIT_NULL on the interface.

An unnumbered IP interface must be created with “peer” address configuration. This peer address is not mandatory to run any dynamic protocol or for static routing. Instead, it is mandatory if any static MPLS services run over this interface.

See also:

- [“Configuring static routes with numbered or unnumbered IP interfaces” on page 5-18](#)
- [“Static MPLS-TP co-routed tunnel support over the unnumbered IP interface” on page 8-2](#)

Step	Action
1	Create the L2 flood domain for the desired VLAN and attach the underlying L2 interfaces. - Create the VLAN and associate it with specific port(s). <pre> vlan create vlan <vlan-id> vlan add vlan <vlan-id> port <port-list> port set port <port> pvid <vlan-id> port set port <port> acceptable-frame-type all </pre> Note: A different VLAN should be used for each IP Interface and physical port combination participating in the VPLS to prevents the creation of flood domains across Layer 3 segments. - If operating in a ring topology, remove VLAN 1,127 from each of the physical ports. <pre> vlan remove vlan 1,127 port <port-list> </pre>

2 Create the IP interface to be used for routing/signaling.

```
interface create ip-interface <ip-interface-name> ip <ip-
address> [vlan <vlan-id>|vs <vs-name>] [mtu <1500-9216>]
[ip-forwarding <on|off>][if-num <NUMBER: 1-4294967295>]
[priority <NUMBER: 0-7>]
```

Note: Be sure to set the MTU size for the IP interface appropriately for the frame size of the data traffic. If set less than the frame size, the data traffic will be dropped.

3 Create the IP loopback interface (optional).

```
interface create loopback <loopback-interface-name> ip
<ip-address>
```

Note: The loopback address is typically used as the router identifier for routing protocols, LDP identifier, and MPLS tunnel peer address.

4 Create the unnumbered IP interface with peer address configuration (optional).

```
interface create ip-interface <ipun1> ip <unnumbered>
donor-interface <lpbk> vlan <8> peer-ip-address <peer
node IP address>
```

Note: The peer node IP address must be the interface address on the other end. For example, it can be a peer node loopback address (if this unnumbered interface is borrowing a loopback IP) or it can be any numbered interface address (if this unnumbered interface is borrowing the numbered interface address).

—end—

Procedure 3-3

Disabling RSTP and MSTP

RSTP and MSTP do not interoperate with MPLS, so ensure these protocols are disabled.

Step	Action

To disable RSTP

1	Disable RSTP: <code>rstp disable</code>

To disable MSTP

2	Disable MSTP: <code>mstp disable</code>

—end—

MPLS overview task flow

This chapter provides an overview of the tasks for configuring MPLS static and dynamic configurations. The general steps for configuring MPLS are:

- 1 Install the MPLS software license. Refer to [“Initial configuration” on page 3-1](#).
- 2 Configure the IP interfaces. Refer to [“Initial configuration” on page 3-1](#).
- 3 Configure the routing protocols. Refer to [“Routing protocol configuration” on page 5-1](#):
 - a. Configure the OSPF routing protocol.
 - b. Configure IS-IS routing protocol.
- 4 Determine whether to use static or dynamic configuration.
- 5 If dynamic, configure the RSVP-TE protocol. Refer to [“RSVP-TE configuration” on page 6-1](#).
- 6 Configure label ranges. Refer to [“Label range configuration” on page 7-1](#).
- 7 Determine whether the switch is an LSR or LER based upon the location of the switch in the network and whether to use IP/MPLS or MPLS-TP tunnels.*
- 8 Configure tunnels.
 - For static tunnels, refer to [“Static tunnel configuration” on page 8-1](#)
 - For dynamic tunnels, refer to [“Dynamic tunnel configuration” on page 9-1](#)
- 9 If dynamic, configure LDP. Refer to [“LDP configuration” on page 11-1](#).
- 10 For LER, configure the virtual circuit or circuits and the virtual switch or switches. Refer to [“Virtual circuit configuration” on page 12-1](#) and [“Virtual switch configuration” on page 13-1](#).
- 11 Configure L2 VPN services. Refer to [“L2 VPN service configuration” on page 14-1](#).
- 12 Configure remote management for MPLS, if desired. Refer to [“Interface configuration” on page 15-1](#)
- 13 Configure VPLS Integrated Routing and Bridging (IRB), if desired. Refer to [“Virtual private LAN service integrated routing and bridging” on page 19-1](#)
- 14 Configure seamless MPLS, if desired. Refer to [“Seamless MPLS configuration” on page 20-1](#)

[Figure 4-1](#) provides an overview of MPLS configuration.

Figure 4-1
MPLS configuration overview





Figure 4-2 provides an overview of static MPLS configuration.

Figure 4-2
MPLS static configuration overview




Refer to [“Static tunnel configuration” on page 8-1](#) for the procedures for

- configuring ingress tunnels
- configuring transit tunnels
- configuring egress tunnels

If the initial configuration needs to be modified, you can modify tunnel attributes. Refer to [“Modifying tunnel attributes” on page 8-33](#).

Figure 4-3 provides an overview of dynamic MPLS configuration.

Figure 4-3
MPLS dynamic configuration overview



Note: You can configure both IS-IS and OSPF. However, to simplify implementation, Ciena recommends using IS-IS or OSPF.

Additional tunnel configuration

Static ingress (unidirectional and bidirectional LSPs) and egress bidirectional LSPs can be routed to an alternate path by changing their next hop or previous hop information. For the procedure, refer to [“Moving co-routed tunnels to a new path” on page 10-10](#).

The network operator can create static co-routed and static or dynamic associated primary and backup tunnel pairs. For the procedures, refer to

- [“Configuring a static co-routed primary tunnel with a dynamic associated tunnel as backup” on page 10-15](#)
- [“Configuring a dynamic associated tunnel as primary and a static co-routed tunnel as backup” on page 10-17](#)

The network operator can create:

- primary dynamic tunnels with explicit paths and secondary dynamic tunnels without explicit paths. For the procedure, refer to [“Configuring a primary tunnel with an explicit path and a backup tunnel without an explicit path” on page 10-19](#).
- primary dynamic tunnels without explicit paths and secondary dynamic tunnels with explicit paths. For the procedure, refer to [“Configuring a primary tunnel without an explicit path and a backup tunnel with an explicit path” on page 10-21](#).
- both primary and secondary dynamic tunnels without explicit paths. For the procedure, refer to [“Configuring a primary tunnel and a backup tunnel without explicit paths” on page 10-23](#).

Routing protocol configuration

Configuring a routing protocol is required in dynamic MPLS deployments and is optional for static MPLS deployments. Routing protocols are IS-IS and OSPF.

Note: You can configure both IS-IS and OSPF. However, to simplify implementation, Ciena recommends using IS-IS or OSPF.

A loopback interface is required in order for higher level protocols, such as RSVP-TE and LDP, to use the route information that OSPF or IS-IS provides.

Static routes can also be configured with numbered or unnumbered IP interfaces.

To configure IS-IS on a point-to-point (P2P) or unnumbered interface, you can enable IP on the interface and then bring it up without assigning a unique IP address to it. You first borrow an IP address already configured on one of the router's other interfaces, then configure IS-IS on it.

To set the interface authentication on a P2P/unnumbered interface, you can configure it for any one level and the interface authentication value is applied to both levels. Similarly, while setting the hello-timer and hello-multiplier values for P2P/unnumbered interface on any one level, the values are reflected for both levels. This implementation is because P2P uses only a single hello packet to represent both level L1 and level L2.

Note: For P2P/unnumbered interfaces, IS-IS internally takes the value of hello-timer, hello-multiplier and interface-authentication from level L1.

In case of P2P/unnumbered interface configuration (hello-timer, hello-multiplier and interface-authentication), the value last configured by the user for level L1 or level L2 is applied to both levels. For example, if the user configured the hello timer for level L1 as 25 and after that configured the hello timer for level L2 as 30, 30 is used for both level L1 and level L2. The last configured value for an IS-IS level for these parameters is used for both levels.

You can configure IS-IS route preference when you:

- Don't want SAOS to choose the OSPF routes as the preferred routes (because the OSPF default administrative distance is smaller than that of IS-IS).
- Have OSPF and IS-IS configured at the same time and then want to switch to IS-IS routes.

SNMP support

This section describes SNMP support for OSPF and IS-IS routing protocols.

OSPF routing protocol

SNMP notification for the OSPF neighbor state change is supported. An SNMP trap is generated for all the OSPF neighbor state changes.

The object identifier (OID) used for this notification is a new Ciena-specific OID for OSPF neighborhood state change, starting with 1.3.6.1.4.1.1271.2.2.104.1.1.

The OSPF neighbor change notification has the following details:

- cienaGlobalSeverity
- cienaGlobalMacAddress
- cienaCesOspfRouterId
- cienaCesOspfNbrIpAddr
- cienaCesOspfNbrAddressLessIndex
- cienaCesOspfNbrRtrId
- cienaCesOspfNbrState

IS-IS routing protocol

An SNMP trap is generated whenever the IS-IS adjacency state is changed to an UP state from a DOWN/INIT state and when the state is changed to a DOWN/INIT from an UP state.

The object identifier (OID) used for this notification is the standard IS-IS OID for adjacency change: 1.3.6.1.2.1.138.0.17. There is no notification for DOWN to INIT and INIT to DOWN state changes.

Procedures

This section provides the following procedures:

- [“Configuring OSPF routing protocol” on page 5-4](#)
- [“Configuring IS-IS routing protocol” on page 5-8](#)

- “Configuring static routes with numbered or unnumbered IP interfaces” on page 5-18
- “Configuring IS-IS on a point-to-point or unnumbered interface” on page 5-19
- “Configuring IS-IS route preference” on page 5-20

Procedure 5-1

Configuring OSPF routing protocol

Configure OSPF.

The default OSPF area is area 0 (IP address 0.0.0.0 with the default type of normal) is automatically created when you first attach an interface to area 0.0.0.0. Once created it cannot be deleted. You can optionally configure an additional non-backbone area enabling the system to perform the functions of an Area Border Router (ABR). With type normal, the continuous backbone area called area 0.0.0.0 is directly connected to every other area and is used for inter-area routing.

Step	Action

To configure OSPF

- 1 Create an OSPF instance:

```
ospf instance create ospf-instance <instance-name>
```

where

ospf-instance is the OSPF instance name.
 <instance-name>
- 2 Modify attributes for the OSPF instance (optional):

```
ospf instance set ospf-instance <ospf-instance>
```

where

ospf-instance is the OSPF instance name.
 <instance-name>

{rfc-1583-compatible <on | off>} turns RFC 1583 compatibility on or off. Default is off to maintain multiple intra AS routes. When AS boundary routers belong to different OSPF areas, and advertise the same external destination, turn on RFC 1583 compatibility.

{router-id <IP address>} is the OSPF router ID. Default is the loopback address.

{spf-delay-interval <MILLISECOND S: 0-4294967295>} is the SPF delay interval in milliseconds. Default is 5000.

3 Create an OSPF area (optional):

```
ospf area create area-ip <ipv4 address> ospf-instance
<instance-name> type <normal|nssa|stub>
```

where

area-ip <ipv4 address> is the IP address of the new non-backbone area.

ospf-instance <instance-name> is the OSPF instance.

type is the type of area.

4 Attach the OSPF instance to an IP interface and OSPF area:

```
ospf interface attach ip-interface <ip-interface> {area
<IpAddress>} {ospf-instance <instance-name>}
```

where

ip-interface <ip-interface> is the IP interface name.

area <IP address> is the OSPF area.

ospf-instance <instance-name> is the OSPF instance name.

5 To advertise an IP interface into OSPF, but without actually running OSPF on that interface, OSPF needs to be run in passive mode. When an IP interface running OSPF passive mode, no OSPF control packets will be originated from the interface. OSPF will advertise the interface as a stub network link through its type 1 link state advertisements (LSAs).

Modify OSPF attributes for an IP interface (optional):

```
ospf interface set ip-interface <ip-interface> {delay-
interval <SECONDS: 1..2147483647>} {bfd-monitor
<enable|disable>}{hello-interval <SECONDS: 1..49999>}
{poll-interval <SECONDS: 1..2147483647>} {priority
<0..255>} {retransmit-delay <SECONDS: 1..3000>}
{transmit-delay <MILLISECONDS: 1..429496799>} {cost-
metric <NUMBER: 0..65535>} {authentication-type <md5 |
text>}{echoless-password}{passive}
```

where

ip-interface <ip-interface> is the IP interface name.

delay-interval <SECONDS: 1..2147483647> is the delay interval of the interface in seconds. Default is 40 seconds.

bfd-monitor <enable|disable> enables or disables bfd-monitor.

where

hello-interval
<SECONDS:
1..49999> is the hello interval of the interface in seconds. Default is 10 seconds.

poll-interval
<SECONDS:
1..2147483647> is the poll interval for the interface in seconds. Default is 120 seconds.

priority <0..255> is the priority used for the interface. Default is 1.

retransmit-delay
<SECONDS:
1..3000> is the retransmit delay interval for the interface in seconds. Default is 5 seconds.

transmit-delay
<MILLISECOND
S:
1..429496799> is the transmit delay interval for the interface in milliseconds. Default is 100 milliseconds.

cost-metric is the cost the interface. Default is 30.

authentication-
type <md5 | text> is the authentication type. Optionally, you can set the authentication type to use MD5 or text authentication.

echoless-
password engages an echoless password collector. Default is blank. Required when the authentication type is set to MD5.

- 6 This step is optional. OSPF area border router (ABR) type 3 LSA filtering provides improved control of route distribution between OSPF areas. OSPF ABR summary type 3 LSA filtering extends the capability of an ABR to filter type 3 LSAs between different OSPF areas. OSPF area filtering filters out specified prefixes to be sent from one area to another area, and allows all other prefixes. This type of area routes filtering can be applied out of a specific OSPF area. It only applies to ABR. OSPF area filtering is supported by the area filter-list command. OSPF ABR type 3 LSA filtering gives the user improved control of route distribution between OSPF areas.

If desired, configure OSPF area route filtering:

```
ospf area add area-ip <IPv4 address> {ospf-instance  
<instance-name>}prefix-filter <IP address>
```

where

area-ip <IPv4
address> is the area IPv4 address.

ospf-instance
<instance-name> is the name of the ospf instance.

prefix-filter is the IP address of the prefix-filter.

To display the configuration

- 7** Display the OSPF instance configuration (optional):
ospf instance show ospf-instance <instance-name>
where
ospf-instance is the OSPF instance name.
<instance-name>
- 8** Display the OSPF area (optional):
ospf area show ospf-instance <instance-name> [area-ip
<area-ip>]
where
ospf-instance is the OSPF instance name.
<instance-name>
area-ip <area-ip> is the area IP address.
statistics displays statistics.
- 9** Display OSPF IP interface information (optional):
ospf interface show ip-interface <ip-interface>
where
ip-interface <ip-interface> is the IP interface name.
statistics shows OSPF interface statistics.

—end—

Procedure 5-2

Configuring IS-IS routing protocol

Configure IS-IS.

IS-IS interfaces always use the IP-IGP-Default BFD profile. Only the transmit and receive intervals can be changed for this profile. For the procedure, refer to [“Managing BFD profiles” on page 17-10](#).

Step	Action

To configure IS-IS global functionality

- Create an IS-IS instance:

```
isis instance create isis-instance <isis-instance> area
<area-id> [level <L1 | L2 | L1L2>]
```

where

isis-instance	is the IS-IS instance name.
<isis-instance>	Note: Only one instance creation is allowed.
logical-id	is the creation index used in configuration.
<NUMBER>]	
area <ISIS area>	is the area identifier, which has a variable length in Hex (AFI.xxxx.xxxx.....).
level < L1 L2 L1L2>	is the routing level.
- Set the authentication algorithm globally to HMAC MD5 for all types of authentication:

```
isis instance set isis-instance isis1 authentication-algo
hmac-md5
```
- Enable route-leak functionality:

```
isis instance set isis-instance <isis-instance> route-
leak <enable | disable>
```

where

isis-instance	is the IS-IS instance name.
<isis-instance>	
route-leak	is enable or disable. Default is enable.
<enable disable>	

4 Enable installation of the default route on receipt of the attached bit:

```
isis instance set isis-instance <isis-instance>
[redistribute | ignore- attach-bit | authentication-algo
 | send-attach-bit | route-leak <enable | disable>]
```

where

isis-instance <isis-instance>	is the IS-IS instance name.
redistribute <enable disable>	To set the redistribution.
ignore- attach-bit <enable disable>	To enable or disable the ignore attach bit processing. Default is disable.
authentication- algo <enable disable>	To enable the HMAC MD5 algorithm IS-IS instance. Default is disable.
send-attach-bit <enable no- overlap-or- redistribute>	To set the attach bit option on an IS-IS instance. Default is no-overlap-or-redistribute.
route-leak <enable disable>	To enable or disable the route leak on an IS-IS instance. Default is enable.

5 Configure the attached bit:

```
isis instance unset isis-instance <isis-instance>
[redistribute | ignore-attach-bit | authentication-algo |
 send-attach-bit | route-leak]
```

where

isis-instance <isis-instance>	is the IS-IS instance name.
redistribute	To unset the redistribution.
ignore- attach-bit <enable disable>	To unset the ignore attach bit processing and change it to the default value.

where

authentication-algo <enable | disable> To unset the HMAC MD5 algorithm IS-IS instance and change it to the default value.

send-attach-bit <enable | no-overlap-or-redistribute> To unset the attach bit option on an IS-IS instance and change it to the default value.

route-leak <enable | disable> To unset the route leak on an IS-IS instance and change it to the default value.

6 Add an area to the IS-IS instance:

```
isis instance add isis-instance <isis-instance> area  
<area-id>
```

where

isis-instance <isis-instance> is the IS-IS instance name.

area <area-id> is the area identifier, which has a variable length in Hex (AFI.xxxx.xxxx.....).

7 Add routes to leak from L2 to L1:

```
isis instance add isis-instance <isis-instance> route-  
leak <l2-to-l1> ip <ip-address/mask>
```

where

isis-instance <isis-instance> is the IS-IS instance name.

route-leak <l2-to-l1> is L2 to L1.

ip <ip-address/mask> is the IP address/mask.

8 Add routes to be summarized from L1 to L2:

```
isis instance add isis-instance <isis-instance> summary-  
route <l1-to-l2> ip <ip-address/mask> [wide-metric  
<metric-value>]
```

where

isis-instance <isis-instance> is the IS-IS instance name.

where

summary- route is L1 to L2.
<l1-to-l2>

ip <ip-address/
mask> is the IP address/mask.

wide-metric is TLVs with link values from 1-16777215.
<metric-value>

- 9** Set the IS-IS domain authentication for level L2 to authenticate sequence number packets (SNPs), including link state packets, CSNPs, and PSNPs (optional):

```
isis domain-authentication set isis-instance <isis-
instance> [authentication-type {md5 | text}] [echoless-
password <password-string>][send-only {yes | no} | snp-
authenticate {yes | no}] [secret <string[256]>]
```

where

isis-instance is the IS-IS instance name.
<isis-instance>

authentication- is the type of authentication. Optionally, you can set the
type {md5 | text} authentication type to use MD5 or text authentication.

echoless- is the interactive password.
password
<password-
string[63]>

send-only {yes | determines whether authentication occurs only upon
no} sending packets. Default is no.

snp-authenticate is the CSNP, PSNP PDU validation.
{yes | no}

secret sets the password using a pre-encrypted string.
<string[256]>

10 Configure the attach bit:

```
isis instance set isis-instance <isis-instance> send-attach-bit <Attached | No-Overlap-Only | Cleared | No-Overlap-or-Redistribute>
```

where

isis-instance is the IS-IS instance name.

<isis-instance>

send-attach-bit Attached indicates the attached bit should always be set.

<Attached | No-Overlap-Only | No-Overlap-Only indicates the attached bit is set only if the IS can reach at least one other area.

Cleared indicates the attached bit should never be set.

Overlap-or-Redistribute> No-Overlap-or-Redistribute indicates that the IS can reach at least one other area (the IS is adjacent with a L2 router whose area addresses do not overlap with the area addresses we know about at L1) OR we have redistribution enabled on the device. This is the default behavior.

11 Change an IS-IS instance level:

```
isis instance set isis-instance <isis-instance> level <L1
| L2 | L1L2>
```

where

isis-instance is the IS-IS instance name.

<isis-instance>

level <L1 | L2 | is the routing level.

L1L2>

12 Set the maximum lifetime and refresh interval to control link state packet generation (optional):

```
isis lsp set isis-instance <isis-instance> [max-lifetime
<NUMBER: 350-65535>] [refresh-interval <NUMBER: 1-65535>]
```

where

isis-instance is the IS-IS instance name.

<isis-instance>

max-lifetime	is the maximum lifetime of an LSP. Default value is 1200
<NUMBER: 350-	seconds.

65535>

refresh-interval is the LSP refresh interval. Default value is 900 seconds.

<NUMBER: 1-

65535>

- 13** Configure SPF calculation settings to control when updates to the link state database occur (optional):

```
isis spf-calculations set isis-instance <isis-instance>
[max-delay <MILLISECONDS>] [threshold-restart-limit
<NUMBER: 1-100>] [threshold-update-restart <NUMBER: 1-
100>] [threshold-update-start <NUMBER: 1-100>]
```

where

isis-instance <isis-instance> is the IS-IS instance name.

max-delay <MILLISECONDS> is the maximum delay (msecs) to start a computation, which determines how long to wait after a database update before updating SPF routing calculations. Default is 5000 milliseconds. When set to 0, the routing calculation occurs immediately following the database update.

threshold-restart-limit <NUMBER: 1-100> is the maximum number of restarts before an in-progress computation run is completed. Default is 10.

threshold-update-restart <NUMBER: 1-100> is the minimum number of changes before the restart of a in-progress computation before interrupting any running SPF routing calculation. Default is -1 meaning that the no interruptions will occur to SPF routing calculations. When set to 0, a database update will cause any running SPF routing calculation to be restarted.

threshold-update-start <NUMBER: 1-100> is the minimum number of changes before the start of computation. Default is -1 meaning that the timing of the SPF routing calculation is determined by the configured calculation delay. When set to 0, any database update will cause an SPF routing calculation to occur.

- 14** Configure IS-IS area authentication to authenticate sequence number packets (SNPs), including link state packets, CSNPs, and PSNPs (optional):

```
isis area-authentication set isis-instance <isis-
instance> [authentication-type {md5 | text}] [echoless-
password <password-string[63]>] [send-only {yes | no} |
snp-authenticate {yes | no}] [secret <string[256]>]
```

where

isis-instance <isis-instance> is the IS-IS instance name.

authentication-type {md5 | text} is the type of authentication. Optionally, you can set the authentication type to use MD5 or text authentication.

echoless-password <password-string[63]> is the interactive password.

where

send-only {yes | no} determines whether authentication occurs only upon sending packets. Default is no.

snp-authenticate {yes | no} is the CSNP, PSNP PDU validation.

secret sets the password using a pre-encrypted string.
<string[256]>

To configure IS-IS interface functionality

- 15** Attach an IP interface to an IS-IS instance:

```
isis interface attach ip-interface <interface-name> isis-  
instance <instance-name> [level <L1 | L2 | L1L2>]
```

where

ip-interface <ip-interface> is the IP interface name.

isis-instance is the IS-IS instance name.
<isis-instance>

level <L1 | L2 | L1L2> is the routing level.

- 16** Configure IS-IS interface authentication at the router level (optional):

```
isis interface-authentication set ip-interface <ip-  
interface> [authentication-type {md5 | text}] [echoless-  
password <password-string[63]>] [send-only {yes | no}]  
[level {L1 | L2 | L1L2}] [secret <String[256]>]
```

where

ip-interface <ip-interface> is the IP interface name.

authentication-type {md5 | text} is the type of authentication. Optionally, you can set the authentication type to use MD5 or text authentication.

echoless-password is the interactive password.
<password-string[63]>

send-only {yes | no} determines whether authentication occurs only upon sending packets. Default is no.

level {L1 | L2 | L1L2} is the routing level.

secret sets the password using a pre-encrypted string.
<string[256]>

17 Modify the IS-IS padded hello for an IP interface (optional):

```
isis interface set ip-interface <interface-name> padded-
hello <yes | no>
```

where

ip-interface <ip-
interface> is the IP interface name.

padded-hello
<yes | no> determines whether to pad hello packets up to the MTU size of the associated IP interface. Default is to use the “padded hello.”

18 Modify IS-IS interface level attributes for an IP interface:

```
isis interface-level set ip-interface <ip-interface>
[level <L1 | L2>] [priority <number>] [hello-timer
<secs>] [hello-multiplier <number>] [wide-metric
<number>] [lsp-interval <msecs>] [csnp-interval <secs>]>
```

where

ip-interface <ip-
interface> is the IP interface name.

[level <L1 | L2>] is the routing level.

[priority
<number>] is the designated router priority. Default is 64. Priority 0 means the interface does not become the designated router.

[hello-timer
<secs>] is the minimum time between successive Hello packets. Range is 1 to 600 seconds. Default is 10 seconds.

[hello-multiplier
<number>] is the hold time for transmitted hello packets. Range is 2 to 200 seconds. Default is 3 seconds.

[wide-metric
<number>] is the cost of the link. Default is 128.

[lsp-interval
<msecs>] is the minimum time between successive link state packets at the level on the circuit. Range is 33 to 65535. Default is 50 milliseconds.

[csnp-interval
<secs>] is the minimum time in seconds between successive CSNPs. Default is 10 seconds.

To redistribute static routes to IS-IS

- 19** Enable redistribution of static routes into IS-IS:

```
isis instance set isis-instance <isis-instance>
redistribute static [metric metric-value] [metric-type
[L1-int | L1-ext> L2-int]
```

where

isis-instance is the IS-IS instance name.

<isis-instance>

[metric metric-value] is the routing level.

metric-type [L1-int | L1-ext> L2-int] is the routing level.

To display the configuration

- 20** Display IS-IS instance protocol information:

```
isis instance show isis-instance <isis-instance>
{attached-interfaces | database | hostname | is-neighbors
{summary | details} | neighbors {summary | details} |
statistics}
```

where

isis-instance is the IS-IS instance name.

<isis-instance>

attached-interfaces is the IS-IS instance with its attachments.

database is the protocol database details.

hostname is the known hostnames of IS neighbors.

is-neighbors is IS neighbors.
{summary | details}

neighbors is the IP address of neighbors.
{summary | details}

statistics is protocol statistics.

- 21** Display IS-IS interface authentication:

```
isis interface-authentication show [ip-interface <name>]
```

where

ip-interface is the IP interface.
<name>

—end—

Example

The following example configures IS-IS with IP BFD monitoring enabled.

```
isis instance create isis-instance MPLS level L1 area 49.0000
isis interface attach ip-interface mpls_3-2 isis-instance MPLS level L1
isis interface set ip-interface mpls_3-2 bfd-monitor enable
```

Procedure 5-3

Configuring static routes with numbered or unnumbered IP interfaces

Specify the egress interface while configuring the static route over any numbered or unnumbered IP interface. This egress interface is an optional parameter and is provided to configure static routes over numbered or unnumbered IP interfaces.

This allows the user to choose the “best” path if there are multiple equal-cost unnumbered paths to the given next-hop while using static routing.

If the user wants to do a static routing, the user must explicitly add the static route for the peer node to provide the reachability information to the local node. This peer route must be configured with the egress interface or else it is considered to be an unresolved route.

See also:

- [“Configuring IP for MPLS interfaces on 39XX/51XX” on page 3-4](#)
- [“Static MPLS-TP co-routed tunnel support over the unnumbered IP interface” on page 8-2](#)
- [“Inserting nodes in an unnumbered network” on page 8-35](#)
- [“Moving MPLS tunnels from numbered to unnumbered IP interfaces” on page 8-39](#)

Step	Action						
1	Specify the egress interface: <pre>ip route add destination <dest-addr> gateway <gw-addr> egress-interface <intf-name></pre> where <table> <tr> <td><dest-addr></td> <td>is the destination IP address</td> </tr> <tr> <td><gw-addr></td> <td>is the gateway IP address</td> </tr> <tr> <td><intf-name></td> <td>is the name of the numbered or unnumbered egress IP interface</td> </tr> </table> —end—	<dest-addr>	is the destination IP address	<gw-addr>	is the gateway IP address	<intf-name>	is the name of the numbered or unnumbered egress IP interface
<dest-addr>	is the destination IP address						
<gw-addr>	is the gateway IP address						
<intf-name>	is the name of the numbered or unnumbered egress IP interface						

Example

Here is sample configuration of a static route with an unnumbered IP interface:

```
interface set ip-interface unnum1 peer-ip-address 2.2.2.2
ip route add destination 2.2.2.2 gateway 2.2.2.2 egress-interface unnum1
```

Procedure 5-4

Configuring IS-IS on a point-to-point or unnumbered interface

Configure IS-IS on a point-to-point (P2P) or unnumbered IP interface.

Step	Action

1	Create an unnumbered IP interface: <pre>interface create ip-interface <unnum-intf-name> ip unnumbered donor-interface <donor-intf-name> ip- forwarding on vlan <vlan-id></pre> where <table> <tr> <td><unnum-intf-name></td> <td>is the name of the unnumbered interface being created.</td> </tr> <tr> <td><donor-intf-name></td> <td>is the name of the interface being borrowed.</td> </tr> <tr> <td><vlan-id></td> <td>is the VLAN ID.</td> </tr> </table>	<unnum-intf-name>	is the name of the unnumbered interface being created.	<donor-intf-name>	is the name of the interface being borrowed.	<vlan-id>	is the VLAN ID.
<unnum-intf-name>	is the name of the unnumbered interface being created.						
<donor-intf-name>	is the name of the interface being borrowed.						
<vlan-id>	is the VLAN ID.						
2	Configure IS-IS on this interface: <pre>isis interface attach ip-interface <unnum-intf-name> isis-instance <isis-instance> level <L1 L2></pre> where <table> <tr> <td>ip-interface <unnum-intf-name></td> <td>is the name of the unnumbered IP interface.</td> </tr> <tr> <td>isis-instance <isis-instance></td> <td>is the name of the IS-IS instance.</td> </tr> <tr> <td>level <L1 L2></td> <td>is the routing level.</td> </tr> </table>	ip-interface <unnum-intf-name>	is the name of the unnumbered IP interface.	isis-instance <isis-instance>	is the name of the IS-IS instance.	level <L1 L2>	is the routing level.
ip-interface <unnum-intf-name>	is the name of the unnumbered IP interface.						
isis-instance <isis-instance>	is the name of the IS-IS instance.						
level <L1 L2>	is the routing level.						

—end—

Procedure 5-5

Configuring IS-IS route preference

Configure IS-IS route preference.

Use this procedure to set the administrative distance of IS-IS to a lower value than OSPF. This results in IS-IS routes being programmed to fault tolerance (FT) by appropriately cleaning up the corresponding OSPF routes.

Note: Administrative distance configurations do *not* work when either Border Gateway Protocol (BGP) or Internal Border Gateway Protocol (iBGP) is being used as an Interior Gateway Protocol (IGP).

Step	Action

1	Configure IS-IS route preference:

	<pre>isis administrative-distance <set unset> [distance <distance-value(1- 255)>]</pre>

where

administrative- distance <set unset>	is the action being taken for the route.

distance <distance- value(1-255)>	is the administrative distance. The default IS-IS administrative distance is 115. The default OSPF administrative distance is 30.

Note: Ciena recommends *not* setting the administrative distance to the same value for both IS-IS and OSPF.

—end—

RSVP-TE configuration

RSVP-TE sets up LSPs in dynamic deployments.

RSVP-TE supports authentication at the IP interface level and at the IP address level and optionally to use MD5 authentication.

For IP address level authentication to work, the parent IP interface authentication must also be configured. Also, authentication configuration between two devices affects the status of MPLS tunnels created by RSVP-TE as shown in [Table 6-1](#).

Table 6-1
Authentication configuration and tunnel status

IP address level configured on both	Interface level configured on both	Message-digest value on both
Yes	No	Same
Yes	No	Different
Yes	Yes	Same at IP address level. Different at interface level.
Yes	Yes	Different at IP address level. Same at interface level.
Yes	One only	Same at IP address level. Default on interface level.
No	Yes	Same
No	Yes	Different
No	One only	Different. Default on IP address and interface levels on non configured device.

Configure RSVP-TE paths to implement RSVP-TE with Explicit Route Object (ERO) for tunnel redundancy or for the following use cases:

- explicit strict route hops can be specified to navigate tunnels through specific hops for cost and congestion avoidance purposes
- explicit route hops could be sparse or network operator can configure every single hop from source to destination. This would be close to configuring static LSP but dynamic case and done only at the head end.
- explicit loose route hops can be specified when the network operator prefers certain locations for tunnel to pass through but may not be sure if TE attributes would be available
- when a tunnel needs to pass through multiple OSPF areas, Area Border Routers should be specified as the explicit route hops
- explicit loose route hops can be specified using either an IP interface address or a loopback address. Explicit strict route hops must be specified using an IP interface address.

RSVP-TE rapid retry

Users can configure a RSVP-TE refresh interval on a node. It is a global command on the node and applies to all the refresh intervals with its RSVP neighbors.

In an auto-backup scenario, after the primary LSP goes down, it takes RSVP more than 157.5 seconds to time out, recalculate the constrained shortest path first (CSPF), and try an alternative primary LSP (new explicit routing object or ERO). The reason is that $R = 30$ seconds in the following RSVP refresh timeout algorithm:

$$\text{Time out} = (K + 0.5) * 1.5 * R = (3 + 0.5) * 1.5 * 30 = 157.5 \text{ seconds}$$

The default is $K = 3$ and $R = 30$ seconds.

To avoid the premature loss of state, L must satisfy $L \geq (K + 0.5) * 1.5 * R$, where K is a small integer. Then in the worst case, $K-1$ successive messages may be lost without the state being deleted. To compute a lifetime L for a collection of states with different R values, such as R_0 , R_1 , and so on, replace R with $\max(R_i)$.

To allow a quicker timeout and to help switching to a new alternative primary path in the auto-backup case, the R value is now configurable with the new CLI `rsvp-te-set refresh-interval` 1 second to 65 seconds.

With $R = 1$ second, and set on all the nodes in the topology, that is, `rsvp-te set refresh-interval 1`, the time out = $(3+0.5) * 1.5 = 5.25$ seconds.

After 5.25 seconds, RSVP begins calculating CSPF and retrying on the next available best path.

This chapter provides the following procedure:

- [“Configuring RSVP-TE” on page 6-4](#)


Procedure 6-1

Configuring RSVP-TE

RSVP-TE sets up LSPs in dynamic deployments. By default RSVP-TE is disabled. The minimum RSVP-TE configuration is to enable RSVP-TE globally and for the IP interface. RSVP-TE will signal over an IP interface when it is enabled.

Figure 6-1 shows the overview for configuring RSVP-TE.

Figure 6-1
MPLS RSVP-TE configuration overview



Step	Action
1	Configure RSVP-TE retry attributes if variation from the default is desired (optional). Disable RSVP-TE. <code>rsvp-te disable</code>

b. Set the retry attributes:

```

rsvp-te set retry-interval <NUMBER: 3...65> retry-
infinite <off>

```

where

retry-interval is the retry time interval in seconds. RSVP-TE attempts to restore tunnels ten times at the retry interval before it gives up. The default setting is 3 seconds.

<NUMBER: 3...65>

retry-infinite <on | off> turns infinite retries for RSVP tunnels on or off. If the retry infinite attribute is turned on, when RSVP tunnels go down because of any reason other than admin = down, RSVP tries to restore the tunnel infinite times. Default is "on".

c. Confirm RSVP-TE retry settings configuration (optional):

```

rsvp-te show [statistics]

```

2 Configure RSVP-TE refresh attributes if variation from the default is desired (optional).**a. Disable RSVP-TE.**

```

rsvp-te disable

```

b. Set the refresh attributes:

```

rsvp-te set refresh-interval <NUMBER: 1...65>

```

where

refresh-interval is the refresh time interval in seconds. The default setting is 30 seconds. See ["RSVP-TE rapid retry" on page 6-2](#) for details.

<NUMBER: 1...65>

c. Confirm RSVP-TE retry settings configuration (optional):

```

rsvp-te show [statistics]

```

3 Enable RSVP-TE globally.

```

rsvp-te enable

```

4 Enable RSVP-TE for the specific IP interface:

```

rsvp-te enable [ip-interface <interface-name>]

```

where

ip-interface is the IP interface.

<interface-name>

5 Configure RSVP-TE IP interface attributes if variation from the default is desired (optional).

a. Set the RSVP-TE attributes for IP interfaces (optional):

```
rsvp-te set ip-interface <name> [advertised-label  
<implicit-null | non-reserved>] [hello-interval  
<NUMBER: 0-30>] [hello-tolerance <NUMBER: 0-10>]  
[authentication-type <md5>] [echoless-password]  
password <Password String>
```

where

ip-interface is the IP interface.
<name>

[advertised-label is the advertised label to support an Implicit Null
<implicit-null |
non-reserved>] Label or a non-reserved label (default).

[hello-interval is the RSVP-TE hello message interval ranging from
<NUMBER: 0-
30>] 0-30 seconds. Default is 0, which is disabled.

[hello-tolerance is the RSVP-TE hello tolerance defines number of
<NUMBER: 0-
10>] hello intervals which may pass without receiving a
 successful Hello message from a partner before the
 Hello session times out. The range is 0-10, and the
 default is 3.

[authentication-
type <md5>] is the authentication type to enable MD5
 authentication. Default is none.

echoless-
password engages an echoless password collector.

password is an 8-30 character password. Required when the
<Password-
String> authentication type is set to MD5. Default is blank.

b. Confirm the RSVP-TE attributes for IP interfaces (optional):

```
rsvp-te interface show
```

6 Configure RSVP-TE IP address authentication (optional).

a. Add an IP address:

```
rsvp-te authentication set peer <ip-address>  
[authentication-type md5] [echoless-password]  
[password <password>]
```

where

peer <ip-
address> is the neighbor IP address.

where

authentication-type <md5>	sets the authentication type to use MD5 authentication. Default is blank.
echoless-password	engages an echoless password collector.
password <password>	is an 8-30 character password. Required when the authentication type is set to MD5. Default is blank.

- b.** Confirm the IP address entry:

```
rsvp-te authentication show
```

7 Configure RSVP-TE paths (optional).

- a.** Create an RSVP-TE path:

```
rsvp-te path create rsvp-path <name>
```

- b.** Set the hops for an RSVP-TE path. For hop-type loose, the parameter ip <ip-address> can be either an IP interface address or a loopback address. For hop-type strict, the parameter ip <ip-address> must be an IP interface address.

```
rsvp-te path set rsvp-path <name> index <number(1-100)> ip <ip-address> [hop-type <strict | loose>]
```

- c.** Confirm configuration of all or specific paths:

```
rsvp-te path show [rsvp-path <name>]
```

—end—

Label range configuration

MPLS label ranges can be modified at any time. However, the chassis must be rebooted in order for the new range to become operational. Also, label ranges between static and dynamic cannot overlap.

Static egress labels can be any valid label between 16-1044479. Static ingress labels can only be from the specified MPLS static label range.

The static pseudowire label range is reserved for MPLS static pseudowires.

The dynamic label range should be same on both ends of a tunnel or virtual circuit. The dynamic label range is shared by virtual circuits and tunnels.

This chapter provides the following procedures:

- [“Configuring label ranges”](#)
- [“Displaying label ranges”](#)

Procedure 7-1

Configuring label ranges

MPLS label ranges can be modified at any time. However, the chassis must be rebooted in order for the new range to become operational. Also, label ranges between static and dynamic cannot overlap.



CAUTION

Tunnels and VCs Could be Removed from Configuration

If there are any tunnels or virtual circuits configured to use labels outside of the new range, they are removed from the configuration upon reboot.

Static incoming labels can only be from the specified MPLS range. The configured static label range determines the valid range for static in labels.

Static outgoing labels can be any valid label between 16-1044479. There are no restrictions on static outgoing labels.

The static pseudowire label range is reserved for MPLS static pseudowires.

The dynamic label range should be same on both ends of a tunnel or virtual circuit. The dynamic label range is shared by virtual circuits and tunnels.

You can configure

- static label ranges. The default minimum label value is 16; the default maximum label value is 4095.
- static pseudowire label ranges. The default minimum label value is 4096; the default maximum label value is 8191.
- dynamic label ranges. The default minimum label value is 8192; the default maximum label value is 1044479.

Table 7-1 describes the rules that apply when configuring label ranges.

Table 7-1
Configuration rules

Scenario	Description
Overlapping label range is not allowed.	As long as the order of different label ranges is maintained as given, you can shrink and expand individual label ranges without overlapping. At least one label each for static-tunnel label range and static-vc label range needs to be reserved, if the user provisions a large dynamic label range.
When the order of static label ranges need to be swapped	Static label ranges cannot be higher than the dynamic label range. If the order in which the individual type of static label ranges appear to need to be swapped, you must first create an 'empty' unused label range in the beginning of the usable label range by shrinking the higher static label range and shifting the lower static label range ahead into the empty label space created. The label ranges can now be swapped by moving the higher static label range to occupy the resultant 'empty' label space created in the beginning of the user label space. See the example on how to swap the order in which static tunnel label range and static pseudowire label range appear.

Step	Action

To configure static label range

- 1 Configure static tunnel label range:

```
mpls static-tunnel-label-range set min-label <NUMBER: 16-1044479> max-label <NUMBER: 16-1044479>
```

To configure static pseudowire label range

- 2 Set the static pseudowire label range:

```
mpls static-vc-label-range set [min-label <NUMBER:16-1044479>] [max-label <NUMBER:16-1044479>]
```

To configure dynamic label range

```

3      Configure dynamic label range:

mpls dynamic-label-range set min-label <NUMBER: 18-
1044479> max-label <NUMBER: 18-1044479>

—end—

```

Examples

The following example sets the static tunnel label range with a minimum label value of 30 and a maximum label value of 1023.

```
> mpls static-tunnel-label-range set min-label 30 max-label 1023
```

The following example sets the static pseudowire label range with a minimum label value of 1024 and a maximum label value of 2047.

```
> mpls static-vc-label-range set min-label 1024 max-label 2047
```

The following example sets the dynamic label range with a minimum label value of 2048 and a maximum label value of 131071.

```
> mpls dynamic-label-range set min-label 2048 max-label 131071
```

The following example shows how to swap the order in which static tunnel label range and static pseudowire label range appear. This example assumes that the given initial label ranges are as follows:

```

> mpls static-tunnel-label-range show
+----- MPLS/GMPLS Static Label Range -----+
| MPLS Admin Min Static-Total Label | 16 |
| MPLS Admin Max Static-Total Label | 8191 |
| MPLS Admin Min Static-Tunnel Label | 16 |
| MPLS Admin Max Static-Tunnel Label | 4095 |
+-----+-----+
| MPLS Oper Min Static-Total Label | 16 |
| MPLS Oper Max Static-Total Label | 8191 |
| MPLS Oper Min Static-Tunnel Label | 16 |
| MPLS Oper Max Static-Tunnel Label | 4095 |
+-----+-----+

> mpls static-vc-label-range show
+----- MPLS/GMPLS Static Label Range -----+
| MPLS Admin Min Static-Total Label | 16 |
| MPLS Admin Max Static-Total Label | 8191 |
| MPLS Admin Min Static-VC Label | 4096 |
| MPLS Admin Max Static-VC Label | 8191 |
+-----+-----+
| MPLS Oper Min Static-Total Label | 16 |
| MPLS Oper Max Static-Total Label | 8191 |
| MPLS Oper Min Static-VC Label | 4096 |
| MPLS Oper Max Static-VC Label | 8191 |
+-----+-----+

> mpls dynamic-label-range show

```

MPLS/GMPLS Dynamic Label Range				
MPLS Admin	Min	Dynamic	Label	8192
MPLS Admin	Max	Dynamic	Label	1044479
MPLS Oper	Min	Dynamic	Label	8192
MPLS Oper	Max	Dynamic	Label	1044479

This example commands follow, along with the output to verify the swapped static-tunnel and static pseudowire label ranges:

Shrink static pseudowire label range:

```
> mpls static-vc-label-range set min-label 8191 max-label 8191
```

Move static tunnel label range into the empty label space to vacate the label space for static pseudowire label range:

```
> mpls static-tunnel-label-range set min-label 4096 max-label 8190
```

Move the static pseudowire label range to the resultant 'empty' label space (16-4095) just created in the beginning of the user label space.

```
> mpls static-vc-label-range set min-label 16 max-label 4095
```

Expand the static-tunnel label range, as required, in the unoccupied label space before dynamic label range:

```
> mpls static-tunnel-label-range set min-label 4096 max-label 8191
```

7-6 Label range configuration

Verify the swapped static-tunnel and static pseudowire label ranges:

```
> mpls static-tunnel-label-range show
+----- MPLS/GMPLS Static Label Range -----+
| MPLS Admin Min Static-Total Label | 16 |
| MPLS Admin Max Static-Total Label | 8191 |
| MPLS Admin Min Static-Tunnel Label | 4096 |
| MPLS Admin Max Static-Tunnel Label | 8191 |
+-----+-----+
| MPLS Oper Min Static-Total Label | 16 |
| MPLS Oper Max Static-Total Label | 8191 |
| MPLS Oper Min Static-Tunnel Label | 4096 |
| MPLS Oper Max Static-Tunnel Label | 8191 |
+-----+-----+

> mpls static-vc-label-range show
+----- MPLS/GMPLS Static Label Range -----+
| MPLS Admin Min Static-Total Label | 16 |
| MPLS Admin Max Static-Total Label | 8191 |
| MPLS Admin Min Static-VC Label | 16 |
| MPLS Admin Max Static-VC Label | 4095 |
+-----+-----+
| MPLS Oper Min Static-Total Label | 16 |
| MPLS Oper Max Static-Total Label | 8191 |
| MPLS Oper Min Static-VC Label | 16 |
| MPLS Oper Max Static-VC Label | 4095 |
+-----+-----+

> mpls dynamic-label-range show
+----- MPLS/GMPLS Dynamic Label Range -----+
| MPLS Admin Min Dynamic Label | 8192 |
| MPLS Admin Max Dynamic Label | 1044479 |
| MPLS Oper Min Dynamic Label | 8192 |
| MPLS Oper Max Dynamic Label | 1044479 |
+-----+-----+
```

Procedure 7-2

Displaying label ranges

You can display

- static label ranges for tunnels
- static label ranges for MPLS pseudowires
- dynamic label ranges

Step	Action

To display static label ranges for tunnels

- 1 Display static label ranges:

```
mpls static-tunnel-label-range show
```

To display static label ranges for MPLS pseudowires

- 2 Display static label ranges for MPLS pseudowires:

```
mpls static-vc-label-range show
```

To display dynamic label ranges

- 3 Display dynamic label ranges:

```
mpls dynamic-label-range show
```

—end—

Example

The following example shows sample output for the mpls static-tunnel-label-range show command.

```
mpls static-tunnel-label-range show
+----- MPLS/GMPLS Static Label Range -----+
| MPLS Admin Min Static-Total Label | 16 |
| MPLS Admin Max Static-Total Label | 8191 |
| MPLS Admin Min Static-Tunnel Label | 16 |
| MPLS Admin Max Static-Tunnel Label | 4095 |
+-----+-----+
| MPLS Oper Min Static-Total Label | 16 |
| MPLS Oper Max Static-Total Label | 8191 |
| MPLS Oper Min Static-Tunnel Label | 16 |
| MPLS Oper Max Static-Tunnel Label | 4095 |
+-----+-----+
```

7-8 Label range configuration

The following example shows sample output for the `mpls static-vc-label-range show` command.

```
mpls static-vc-label-range show
```

+----- MPLS/GMPLS Static Label Range -----+					
MPLS Admin Min	Static-Total	Label	16		
MPLS Admin Max	Static-Total	Label	8191		
MPLS Admin Min	Static-VC	Label	4096		
MPLS Admin Max	Static-VC	Label	8191		
+-----+					
MPLS Oper Min	Static-Total	Label	16		
MPLS Oper Max	Static-Total	Label	8191		
MPLS Oper Min	Static-VC	Label	4096		
MPLS Oper Max	Static-VC	Label	8191		
+-----+					

The following example shows sample output for the `mpls dynamic-label-range show` command.

```
mpls dynamic-label-range show
```

+----- MPLS/GMPLS Dynamic Label Range -----+				
MPLS Admin Min	Dynamic	Label	8192	
MPLS Admin Max	Dynamic	Label	1044479	
MPLS Oper Min	Dynamic	Label	8192	
MPLS Oper Max	Dynamic	Label	1044479	
+-----+				

Static tunnel configuration

This chapter explains how to configure static tunnels on the 39XX/51XX Service Delivery, Aggregation and Virtualization switches. Static tunnels are supported for TP and IP/MPLS deployments.

Static TP tunnels are:

- static ingress co-routed (static-ingress-corout)
- static egress co-routed (static-egress-corout)
- static transit co-routed (static-transit-corout)
- static ingress unidirectional (static-ingress-unidir)
- static egress unidirectional (static-egress-unidir)
- static transit unidirectional (static-transit-unidir)
- static associated bidirectional (static-ingress-assoc)

Static IP/MPLS tunnels are:

- static ingress (static-ingress)
- static egress (static-egress)
- static transit (static-transit)
- associated bidirectional (bidir-ingress-assoc)

Next-hop diversity

It is recommended that a backup tunnel takes a diverse path from the primary tunnel it is protecting so that failure on the intersecting node or link does not cause the primary and backup tunnel to fail simultaneously. When the network topology prevents this, there is no backup tunnel or a backup tunnel that may share portions of a network topology. The second option provides protection if the network nodes/links that fail are not shared between the primary and backup.

When configuring the ingress unidirectional and ingress and egress co-routed bidirectional static tunnels, users can select next (ingress) or previous (egress) hop be not accessible through the same link as the primary LSP. This

is expressed by “recovery-(p/n)hop-disjoint <none | link>” when configuring primary LSP. The default is link which forces you to provide a link diverse path. Option “none” lets you share the link between primary and backup.

Static MPLS-TP co-routed tunnel support over the unnumbered IP interface

Currently the deployment of any SAOS 6.x node in the customer’s access or aggregation ring needs an IP address per interface for the provisioning of static MPLS-TP co-routed tunnels. All the static MPLS-TP co-routed tunnels are configured with next-hop IP address as the IP address of the next node in the ring.

When any new node is inserted in a ring, the next-hop IP address needs to be changed in the existing MPLS-TP static co-routed tunnels on the directly connected nodes.

SAOS supports the creation of static MPLS-TP co-routed tunnels over the unnumbered IP interface.

Insertion of new nodes in existing access or aggregation rings is painless because static MPLS-TP tunnels can be configured with the unnumbered IP interface, rather than the next/previous hop IP address.

Related procedures are:

- [“Inserting nodes in an unnumbered network” on page 8-35](#)
- [“Moving MPLS tunnels from numbered to unnumbered IP interfaces” on page 8-39](#)

See also:

- [“Configuring IP for MPLS interfaces on 39XX/51XX” on page 3-4](#)
- [“Configuring static routes with numbered or unnumbered IP interfaces” on page 5-18](#)

Attributes for static tunnel configuration are listed in [Table 8-1](#).

Table 8-1
Attributes for static tunnel configuration

Attribute	Description	Applicable Tunnel Type
ais-monitor <enable disable>	Enables or disables AIS based fault monitoring.	static-ingress-corout static-egress-corout static-transit-corout static-transit-unidir static-ingress-assoc

Attribute	Description	Applicable Tunnel Type
ais-profile <MplsAisProfiles>	Specifies the AIS profile to be used for AIS monitoring.	static-ingress-corout static-egress-corout static-transit-corout static-transit-unidir static-ingress-assoc
backup-tunnel <Primary Ingress Tunnel>	Specifies the ingress tunnel to protect the primary tunnel.	static-ingress static-ingress-corout static-egress-corout static-ingress-unidir
bfd-monitor <enable disable>	Enables or disables BFD based fault monitoring.	bidir-ingress-assoc static-ingress-corout static-egress-corout static-ingress-assoc
bfd-profile <BfdProfile>	Specifies the BFD profile to be used for BFD monitoring.	bidir-ingress-assoc static-ingress-corout static-egress-corout static-ingress-assoc
cos-profile <CoS Profile Name>	Selects the tunnel CoS profile.	static-ingress static-transit static-ingress-corout static-egress-corout static-transit-corout static-ingress-unidir static-transit-unidir

8-4 Static tunnel configuration

Attribute	Description	Applicable Tunnel Type
dest-ip <IP Address>}	Sets the destination IP address for the tunnel. This attribute is optional; however, if you enter it, you must also enter • lsp-id • src-tunnel-id These attributes are required to verify the connection for LSP-BFD and LSP ping/traceroute interoperability. Identical information must be configured at the other end.	static-ingress static-transit static-ingress-corout static-transit-corout static-ingress-unidir static-transit-unidir
dest-tunnel-id <#1..8192>	Specifies the ingress tp-tunnel destination Tunnel-ID.	static-ingress-corout static-egress-corout static-transit-corout static-ingress-unidir static-egress-unidir static-transit-unidir
fixed-ttl <NUMBER: 1..255> fixed TTL value	Specifies the fixed TTL value.	static-ingress static-transit
forward-in-label <#16..1044479>	Specifies the forward decap tp-tunnel incoming label.	static-egress-corout static-transit-corout static-egress-unidir static-transit-unidir
forward-out-label <#16..1044479>	Specifies the forward encap tp-tunnel outgoing label.	static-ingress-corout static-transit-corout static-ingress-unidir static-transit-unidir
forward-tunnel <MplsIngressTunl>	Specifies the forward tunnel.	bidir-ingress-assoc static-ingress-assoc
in-label <NUMBER: 16-1044479>	Sets the tunnel in label.	static-egress static-transit

Attribute	Description	Applicable Tunnel Type
lsp-id <#1..65535>	Specifies the ingress tp-tunnel LSP-ID. This attribute is optional; however, if you enter it, you must also enter • dest-tunnel-id • src-tunnel-id These attributes are required to verify the connection for LSP-BFD and LSP ping/traceroute interoperability. Identical information must be configured at the other end.	static-ingress-corout static-egress-corout static-transit-corout static-ingress-unidir static-egress-unidir static-transit-unidir
next-hop-ifNum <NUMBER: 1-4294967295>	Associates a next hop that is not IP-specific so that non-IP LSP ping and traceroute can be run on the tunnel.	static-ingress-corout static-transit-corout
next-hop-ip <IP Address>}	Specifies the IP address of the next hop device. Note: The applicable tunnel types have a dependency on the local interface for the next hop device. Ensure that the local interface is created on the next hop device before creating the tunnel.	static-ingress static-transit static-ingress-corout static-transit-corout static-ingress-unidir static-transit-unidir
next-hop-local interface <unnumbered-if-name>	Specifies the name of the unnumbered IP interface for the next hop.	static-ingress-corout static-transit-corout
out-label <NUMBER: 16-1044479>	Sets the tunnel out label.	static-ingress static-transit
prev-hop-ifNum	Associates a previous hop that is not IP-specific so that non-IP LSP ping and traceroute can be run on the tunnel.	static-egress-corout static-transit-corout

8-6 Static tunnel configuration

Attribute	Description	Applicable Tunnel Type
prev-hop-ip <IpAddress>	Specifies the tp-tunnel previous hop IP address. Note: The applicable tunnel types have a dependency on the local interface for the previous hop device. Ensure that the local interface is created on the previous hop device before creating the tunnel.	static-egress-corout static-transit-corout static-egress-unidir static-transit-unidir
prev-hop-local-interface <unnumbered-if-name>	Specifies the name of the unnumbered IP interface for the previous hop.	static-egress-corout static-transit-corout
queue-group-instance <NUMBER: 1..65534>	Specifies an instance of a queue group profile that is attached to an egress port. The three-tuple of {<queue-group-profile>, <queue-group-instance>, <egress-port>} is required to be unique for each system.	static-ingress static-ingress-corout static-ingress-unidir static-egress-corout
queue-group <qgrp-name>	Specifies the custom queue group profile.	static-ingress static-ingress-corout static-ingress-unidir static-egress-corout
recovery-nhop-disjoint <link none>	Sets the next (for ingress) or previous (for egress) hop diversity option. • ink: provide a link diverse path • none: link between primary and backup is fate-shared The default value is link.	static-ingress static-ingress-corout static-ingress-unidir
reverse-in-label <#16..1044479>	Specifies the reverse decap tp-tunnel incoming label.	static-ingress-corout static-transit-corout
reverse-out-label <#16..1044479>	Specifies the reverse encaps tp-tunnel outgoing label.	static-egress-corout static-transit-corout
reverse-static-tunnel <MplsEgressTunl>	Specifies the static reverse tunnel.	bidir-ingress-assoc static-ingress-assoc

Attribute	Description	Applicable Tunnel Type
reverse-dyntun-name <String>	Specifies the name of the dynamic reverse tunnel.	bidir-ingress-assoc
reversion-hold-time <NUMBER: 0-3600>	Sets the amount of time in seconds to wait before switching from the backup to the primary tunnel after the fault on the primary is cleared (tunnel reversion). Default is 300 seconds. Applicable when tunnel reversion is turned on.	static-ingress static-ingress-corout static-egress-corout static-ingress-unidir
src-ip <IP Address>	Specifies the tunnel source IP address.	static-egress static-transit static-ingress-corout static-egress-corout static-transit-corout static-egress-unidir static-transit-unidir
src-tunnel-id <#1..8192>	Specifies the ingress tp-tunnel source Tunnel-ID. This attribute is optional; however, if you enter it, you must also enter • dest-tunnel-id • lsp-id These attributes are required to verify the connection for LSP-BFD and LSP ping/traceroute interoperability. Identical information must be configured at the other end.	static-ingress-corout static-egress-corout static-transit-corout static-ingress-unidir static-egress-unidir static-transit-unidir
static-egress <static-egress>	Specifies the static egress tunnel name.	static-egress
static-egress-corout <static-egress-corout>	Specifies the egress corouted tunnel name.	static-egress-corout
static-egress-unidir <static-egress-unidir>	Specifies the egress unidirectional tunnel name.	static-egress-unidir
static-ingress <static-ingress>	Specifies the static ingress tunnel name.	static-ingress

Attribute	Description	Applicable Tunnel Type
static-ingress-corout <static-ingress-corout>	Specifies the static ingress corouted tunnel name.	static-ingress-corout
static-ingress-unidir <static-ingress-unidir>	Specifies the static ingress unidirectional tunnel name.	static-ingress-unidir
static-ingress-assoc <static-ingress-accoc>	Specifies the static ingress associated tunnel name.	static-ingress-assoc
static-transit <static-transit>	Specifies the static transit tunnel name.	static-transit
static-transit-corout <static-transit-corout>	Specifies the static transit corouted tunnel name.	static-transit-corout
static-transit-unidir <static-transit-unidir>	Specifies the static transit unidirectional tunnel name.	static-transit-unidir
ttl-policy <string>	Specifies the TTL policy.	static-ingress static-transit
tunnel-reversion <on/off>	Indicates whether to switch from the backup to the primary tunnel after the fault on the primary is cleared.	static-ingress static-ingress-corout static-egress-corout static-ingress-unidir

Procedures

This chapter provides the following procedures:

- [“Configuring static TE tunnels” on page 8-10](#)
- [“Configuring static transit unidirectional TP tunnels” on page 8-12](#)
- [“Configuring static unidirectional ingress TP tunnels” on page 8-13](#)
- [“Configuring static unidirectional egress TP tunnels” on page 8-15](#)
- [“Configuring static bidirectional ingress-associated TE tunnels” on page 8-16](#)
- [“Configuring co-routed TP tunnels” on page 8-17](#)
- [“Displaying MPLS TE-tunnel information” on page 8-24](#)
- [“Displaying TP tunnel information” on page 8-28](#)
- [“Modifying tunnel attributes” on page 8-33](#)
- [“Inserting nodes in an unnumbered network” on page 8-35](#)
- [“Moving MPLS tunnels from numbered to unnumbered IP interfaces” on page 8-39](#)

MPLS tunnels can be protected by configuring backup tunnels at the head-end LER. The backup tunnel provides protection end-to-end for the primary tunnel by means of a set of LSR hops that are primary-path-negated. This ensures that failure on primary path does not translate to simultaneous failure on the backup path.

Tunnel FEC for static LSP

When configuring static LSP, you are encouraged to enter the following additional parameters, as stated in RFC 6426:

- Destination tunnel ID
- Source Tunnel ID
- LSP ID

These parameters are optional, however if you enter one, you have to enter all parameters. These parameters are necessary to verify the connection for LSP-BFD and LSP ping/traceroute interoperability. Identical information must be configured at the other end. The dynamic tunnels exchange this information in signaling which is why it is only needed for static tunnel configuration.

Procedure 8-1

Configuring static TE tunnels

Configure TE tunnels in TE deployments. A static TE deployment requires tunnel creation. The type of tunnel depends on the role of the MPLS switch.

For ingress and egress LERs, create both an ingress and egress tunnel.

For an LSR, create a transit tunnel.

Note: The value for in-label must be in the configured static MPLS label range.

You can configure

- a static ingress tunnel, that is, the ingress LER of a static tunnel
- a static egress tunnel, that is, the egress LER of a static tunnel
- a static transit tunnel, that is, the transit LSR of a static tunnel
- an associated bidirectional ingress TE tunnel

Note: The prev-hop-ip attribute is required for static egress tunnels and static transit tunnels.

Unnumbered IP interface attributes are used when configuring MPLS tunnels that are connected to a neighbor using non-IP unnumbered interfaces. Instead of using the next-hop-IP attribute, static ARP configuration can be used to specify the MAC address of the next hop. For the procedure, refer to [“Adding a static ARP entry” on page 15-8](#).

Step	Action

To configure a static ingress tunnel

- 1 Create a static ingress tunnel:

```
mpls tunnel create static-ingress <static-ingress>
[dest-ip <IP address>] [next-hop-ip <IP address>] {out-
label <NUMBER: 16-1044479>} [cos-profile <MPLS Tunnel COS
Profile>] ttl-policy <String> fixed-ttl <NUMBER: 1..255>
[recovery-nhop-disjoint <none | link>][backup-tunnel
<MPLS ingress primary tunnel>] [reversion-hold-time
<SECONDS: 0-3600>] [tunnel-reversion <on | off>] [queue-
group <qgrp-name>] [queue-group-instance <NUMBER:
1..65534>]
```

To configure a static egress tunnel

- 2 Create a static egress tunnel:

```
mpls tunnel create static-egress <static-egress> [src-ip
<IP address>] {in-label <NUMBER: 16-1044479>} {prev-hop-
ip <IP Address>}
```

To configure a static transit tunnel

- 3 Create a static transit tunnel:

```
mpls tunnel create static-transit <static-transit>
{dest-ip <IP address>} {src-ip <IP address>} {next-hop-ip
<IP address>} {in-label <NUMBER: 16-1044479>} {out-label
<NUMBER: 16-1044479>} {out-explicit-null-label} {out-no-
label} [cos-profile <MPLS Tunnel COS Profile>] ttl-policy
<String> fixed-ttl <NUMBER: 1..255> {prev-hop-ip <IP
Address>}
```

To configure an associated bidirectional ingress TE tunnel

- 4 Create an associated bidirectional ingress TE-Tunnel:

```
mpls tunnel create bidir-ingress-assoc <bidir-ingress-
assoc> {forward-tunnel <MPLS ingress tunnel>} {reverse-
static-tunnel <MPLS egress tunnel>} {reverse-dyntun-name
<String>} [bfd-monitor <enable | disable>] [bfd-profile
<MPLS BFD Profile List>]
```

—end—

Example

The following example configures an MPLS-TE static ingress tunnel.

```
mpls tunnel create static-ingress st-1.1.1.1-A dest-ip 1.1.1.1 next-hop-ip
11.11.11.50 out-label 300
```

The following example configures an MPLS-TE static egress tunnel.

```
mpls tunnel create static-egress st-frm-1.1.1.1 src-ip 1.1.1.1 in-label 400
prev-hop-ip 192.168.1.2
```

The following example configures an MPLS-TE static transit tunnel.

```
mpls tunnel create static-transit frm-10.10.10.10-to-1.1.1.1 dest-ip 1.1.1.1
src-ip 10.10.10.10 next-hop-ip 42.1.1.15 in-label 300 out-label 301
mpls tunnel create static-transit frm-1.1.1.1-to-10.10.10.10 dest-ip
10.10.10.10 src-ip 1.1.1.1 next-hop-ip 11.11.11.100 in-label 401 out-label 400
prev-hop-ip 192.168.1.2
```

The following example configures an MPLS-TE associated bidirectional tunnel.

```
mpls tunnel create bidir-ingress-assoc st-ing-assoc-AP forward-tunnel st-
1.1.1.1-A reverse-static-tunnel st-frm-1.1.1.1
```

Procedure 8-2

Configuring static transit unidirectional TP tunnels

Configure the static transit unidirectional TP tunnels.

Note: Ciena recommends using the optional prev-hop-ip attribute for SAOS 6.14. The prev-hop-ip attribute is used to identify the local interface into which the packet is traversing.

Step	Action
1	Create a static transit unidirectional TP Tunnel: <pre>gmpls tp-tunnel create static-transit-unidir <static-transit-unidir> {dest-ip <IP address>} {src-ip <IP address>} {next-hop-ip <IP address>} {forward-out-label <NUMBER: 16-1044479>} {forward-in-label <NUMBER: 16-1044479>} [cos-profile <MPLS Tunnel COS Profile>] {ttl-policy <fixed decrement>} [fixed-ttl <NUMBER: 1-255>] [prev-hop-ip <IP address>] [ais-monitor <enable disable>] [ais-profile <AIS Profile List>]</pre>
2	Confirm the configuration: <pre>gmpls tp-tunnel show static-transit-unidir <static-transit-unidir></pre>

—end—

Example

```
gmpls tp-tunnel create static-transit-unidir asoc-frm-10.10.10.10-to-1.1.1.1
dest-ip 1.1.1.1 src-ip 10.10.10.10 next-hop-ip 42.1.1.15 forward-in-label
1003 forward-out-label 1001 prev-hop-ip 43.1.1.10
gmpls tp-tunnel create static-transit-unidir asoc-frm-1.1.1.1-to-10.10.10.10
dest-ip 10.10.10.10 src-ip 1.1.1.1 next-hop-ip 11.11.11.100 forward-in-label
1000 forward-out-label 1002 prev-hop-ip 43.1.1.10
```

Procedure 8-3

Configuring static unidirectional ingress TP tunnels

Configure static ingress unidirectional TP tunnels.

Step	Action

1 Configure the primary active path:

```
gmpls tp-tunnel create static-ingress-unidir <static-
ingress-unidir> dest-ip <IpAddress> [src-tunnel-id
<NUMBER>] [forward-lsp-id <NUMBER>] dest-tunnel-id
<NUMBER>] next-hop-ip <IpAddress> forward-out-label
<#16..1044479> [cos-profile <MplsTunlCosProfile> [ttl-
policy <fixed>] [recovery-nhop-disjoint <none|link>]
fixed-ttl <#1..255>
```

2 Configure the backup path:

```
gmpls tp-tunnel create create static-ingress-unidir
<backup tunnel name> dest-ip <dest ip address> next-hop-
ip <next hop ip address> forward-out label <out-label>
backup-tunnel <active tunnel name>
```

3 Confirm the configuration:

```
gmpls tp-tunnel show static-ingress-unidir
<MplsIngressStaticTpUniDirTunl>
```

—end—

Example

The following example configures a static unidirectional ingress TP tunnel.

```
gmpls tp-tunnel create static-ingress-unidir st-ing-u-A dest-ip 1.1.1.1 next-
hop-ip 11.11.11.50 forward-out-label 1003
gmpls tp-tunnel show static-ingress-unidir st-ing-u-A
```

-----GMPLS INGRESS TP-TUNNEL DETAILS-----	
Parameter	Value
Tunnel Name	st-ing-u-A
Tunnel Index	4
Tunnel Type	Static
Direction	Unidir
Nodal Role	Ingress
Destination IP Address	1.1.1.1
Source IP Address	10.10.10.10
Next-Hop IP Address	11.11.11.50
Admin State	Enabled
Oper State	Enabled
Forward Out-Label	1003
Forward Tunnel Group Index	32772
Forward Protection Role	Primary
Forward Protection State	Active
Forward Backup Tunnel Name	None Present

8-14 Static tunnel configuration

Forward Tunnel Reversion	On
Forward Reversion Hold-Time	30
Forward CoS Profile Name	DefaultTunlCoSProfile
Forward CoS Profile Index	1
TTL Policy	fixed
Fixed TTL	255

The following example configures a static unidirectional ingress TP tunnel with backup protection.

```
gmpls tp-tunnel create static-ingress-unidir sta-ing-u-tp-P dest-ip 1.1.1.1
next-hop-ip 5.1.1.1 forward-out-label 2003
```

```
gmpls tp-tunnel create static-ingress-unidir sta-ing-u-tp-B dest-ip 1.1.1.1
next-hop-ip 4.1.1.1 forward-out-label 2040 backup-tunnel sta-ing-u-tp-P
```

```
gmpls tp-tunnel show
```

Flags : P -> Primary B -> Backup
 A -> Active S -> Standby
 L -> Lone Member R -> Recovery Group Member
 E -> Tunnel uses Explicit Path

-----GMPLS INGRESS TP-TUNNEL TABLE-----							
Type	Tunnel Index	Tunnel Name	Destination IP	Out Label	Admin State	Oper State	Flags
Static	11	sta-ing-u-tp-P	1.1.1.1	2003	ENA	ENA	P R A
Static	12	sta-ing-u-tp-B	1.1.1.1	2040	ENA	ENA	B R S

Procedure 8-4

Configuring static unidirectional egress TP tunnels

Configure static unidirectional egress TP tunnels.

Note: Ciena recommends using the optional prev-hop-ip attribute. The prev-hop-ip attribute is used to identify the local interface into which the packet is traversing.

Step	Action

1	Create a static egress unidirectional TP Tunnel:

	<pre>gmpls tp-tunnel create static-egress-unidir <tp-tunnel- name> src-ip <IpAddress> forward-in-label <#16..1044479> prev-hop-ip <IpAddress> [src-tunnel-id <NUMBER>] [forward-lsp-id <NUMBER>][dest-tunnel-id <NUMBER>] [recovery-phop-disjoint <none link>]</pre>

2	Confirm the configuration:

	<pre>gmpls tp-tunnel show static-egress-unidir <MplsEgressStaticTpUniDirTunl></pre>

—end—

Example

```
gmpls tp-tunnel create static-egress-unidir st-egr-u-A src-ip 1.1.1.1
forward-in-label 1002 prev-hop-ip 192.168.1.2 logic-id 10 src-tunnel-id 10
forward-lsp-id 1 dest-tunnel-id 10
```

```
gmpls tp-tunnel create static-egress-unidir st-egr-u-B src-ip 2.2.2.2
forward-in-label 1050 prev-hop-ip 192.168.2.2 logic-id 11 src-tunnel-id 11
forward-lsp-id 1 dest-tunnel-id 11
```

Procedure 8-5

Configuring static bidirectional ingress-associated TE tunnels

Configure both ends of static bidirectional ingress-associated TE tunnels.

Step	Action
1	Create a bidirectional ingress associated TE-Tunnel: <pre>mpls tunnel create bidir-ingress-assoc <bidir-ingress- assoc> {forward-tunnel <MPLS ingress tunnel>} {reverse- static-tunnel <MPLS egress tunnel>} {reverse-dyntun-name <String>} [reverse-dyntun-srcip <IP address>] [bfd- monitor <enable disable>] [bfd-profile <MPLS BFD Profile List>] [ais-monitor <enable disable>] [ais- profile <AIS Profile List>]</pre> —end—

Example

TE-Associated tunnel creation of st-ing-associ-AP with a forward-tunnel (ingress) of st-1.1.1.1-A and the reverse-static-tunnel of st-frm-1.1.1.1 (egress) is shown in the examples for [“Configuring static TE tunnels” on page 8-10](#).

```
mpls tunnel create bidir-ingress-assoc st-ing-associ-AP forward-tunnel st-
1.1.1.1-A reverse-static-tunnel st-frm-1.1.1.1
```

Procedure 8-6

Configuring co-routed TP tunnels

MPLS-TP co-routed tunnels are used for intra-metro TDM PWs.

Co-routed TP tunnels can be:

- static ingress
- static egress
- static transit

Table 8-2 lists the default values for static ingress and egress co-routed TP tunnels.

Table 8-2
Default values for static co-routed TP tunnels

Attribute	Default value
ttl-policy	fixed
fixed-ttl	255
reversion-hold-time	30
tunnel-reversion	on
bfd-monitor	disable
bfd-profile	LSP_BFD_DEF_PROF
ais-monitor	disable
ais-profile	AIS_DEF_PROF

To configure co-routed TP tunnels to be compatible with LSP ping and LSP traceroute in a non-IP environment, use the next-hop-ifNum and prev-hop-ifNum attributes.

Non-IP environments are supported through unnumbered interface support for MPLS-TP LSP configuration. In addition to the neighbor's interface IP address, the following information must be configured:

- next-hop-ifNum when configuring ingress and intermediate co-routed LSP when the next-hop neighbor is the same
- prev-hop-ifNum when configuring egress and intermediate co-routed LSP when the next-hop neighbor is the same

Unnumbered IP interface attributes are used when configuring MPLS tunnels that are connected to a neighbor using non-IP unnumbered interfaces. Instead of using the next-hop-IP attribute, static ARP configuration can be used to specify the MAC address of the next hop. For the procedure, refer to [“Adding a static ARP entry” on page 15-8](#).

Note: The value for forward-in-label of the egress tunnel must be in the configured static MPLS tunnel label range.

Step	Action

To configure a static ingress co-routed TP tunnel

- 1 Create a static ingress co-routed TP-Tunnel:

```
gmpls tp-tunnel create static-ingress-corout <static-ingress-corout> [dest-tunnel-id <NUMBER: 1-65535>] {lsp-id <NUMBER: 1-65535>} [src-tunnel-id <NUMBER>]{dest-ip <IP address>} [next-hop-ip <IP address> | next-hop-local-interface <unnumbered-if-name>] {next-hop-ifnum <NUMBER: 1-4294967295>}{forward-out-label <NUMBER: 16-1044479>} {reverse-in-label <NUMBER: 16-1044479> [cos-profile <MPLS Tunnel COS Profile>] [recovery-nhop-disjoint <none|link>] [backup-tunnel <MPLS ingress primary tp corout tunnel>] [reversion-hold-time <NUMBER: 0-3600>] [tunnel-reversion <on | off>] [bfd-monitor <enable | disable>] [bfd-profile <MPLS BFD Profile List>] [ais-monitor <enable | disable>] [ais-profile <AIS Profile List>] [queue-group <qgrp-name>] [queue-group-instance <NUMBER 1-65534>]
```
- 2 Confirm static ingress co-routed TP-Tunnel creation:

```
gmpls tp-tunnel show static-ingress-corout <static-ingress-corout>
```

To configure a static egress co-routed TP tunnel

- 3 Create a static egress co-routed TP-Tunnel:

```
gmpls tp-tunnel create static-egress-corout <static-egress-corout> [dest-tunnel-id <NUMBER>] {lsp-id <NUMBER: 1-65535>} [src-tunnel-id <NUMBER: 1-65535>] {src-ip <IP address>} [prev-hop-ip <IP address>] | prev-hop-local-interface <unnumbered-if-name>] {prev-hop-ifnum <NUMBER: 1-4294967295>} forward-in-label <NUMBER: 16-1044479> reverse-out-label <NUMBER: 16-1044479> [cos-profile <MPLS Tunnel COS Profile>] [reversion-hold-time <NUMBER: 0-3600>] [tunnel-reversion <on | off>] [recovery-phop-disjoint <none|link>] [backup-tunnel <MPLS static egress primary tp-tunnel>] [bfd-monitor <enable | disable>]
```

```
[bfd-profile <MPLS BFD Profile List>][ais-monitor <enable
| disable>] [ais-profile <AIS Profile List>] [queue-group
<qgrp-name>] [queue-group-instance <NUMBER 1-65534>]
```

4 Confirm static egress co-routed TP-Tunnel creation:

```
gmpls tp-tunnel show static-egress-corout <static-egress-
corout>
```

To configure a static transit co-routed TP tunnel

5 Create a static transit co-routed TP-Tunnel:

```
gmpls tp-tunnel create static-transit-corout <static-
transit-corout> [dest-tunnel-id <NUMBER: 1-65535>] {lsp-
id <NUMBER: 1-65535>} {src-tunnel-id <NUMBER: 1-65535>}
dest-ip <IP address> src-ip <IP address> next-hop-ip <IP
address> {next-hop-ifnum <if-num> | next-hop-local
interface <unnumbered-if-name>} [prev-hop-ip <IP address>
| [prev-hop-local-interface <unnumbered-if-name>] prev-
hop-ifnum <NUMBER: 1-4294967295>] forward-in-label
<NUMBER: 16-1044479> forward-out-label <NUMBER: 16-
1044479> reverse-in-label <NUMBER: 16-1044479> reverse-
out-label <NUMBER: 16-1044479> [cos-profile <MPLS Tunnel
COS Profile>] [ais-monitor <enable | disable>] [ais-
profile <MPLS AIS Profiles>]
```

6 Confirm static transit co-routed TP-Tunnel creation:

```
gmpls tp-tunnel show static-transit-corout <static-
transit-corout>
```

—end—

Example

The following example creates a GMPLS static co-routed ingress TP-tunnel.

```
gmpls tp-tunnel create static-ingress-corout icor-to-1.1.1.1 dest-ip 1.1.1.1
next-hop-ip 5.1.1.1 forward-out-label 2003 reverse-in-label 2002
```

The following sample output confirms the configuration.

```
gmpls tp-tunnel show static-ingress-corout icor-to-1.1.1.1
+-----GMPLS INGRESS TP-TUNNEL DETAILS-----+
| Parameter                                     | Value                                     |
+-----+-----+
| Tunnel Name                                 | icor-to-1.1.1.1                         |
| Tunnel Index                               | 6                                       |
| Tunnel Type                               | Static                                  |
| Direction                                 | Bidir                                  |
| Nodal Role                                | Ingress                                 |
| Destination IP Address                     | 1.1.1.1                                 |
| Source IP Address                          | 10.10.10.10                             |
| Next-Hop IP Address                        | 5.1.1.1                                 |
| Admin State                               | Enabled                                 |
| Oper State                                | Enabled                                 |
| Forward Out-Label                          | 2003                                    |
| Reverse In-Label                           | 2002                                    |
| Forward Tunnel Group Index                 | 32774                                   |
| Forward Protection Role                     | Primary                                 |
```

8-20 Static tunnel configuration

Forward Protection State	Active
Forward Backup Tunnel Name	None Present
Forward Tunnel Reversion	On
Forward Reversion Hold-Time	30
Forward CoS Profile Name	DefaultTunlCoSProfile
Forward CoS Profile Index	1
TTL Policy	fixed
Fixed TTL	255
BFD Monitoring	Disabled
BFD Profile ID	2
BFD Profile Name	Active-LSP
BFD Session ID	3
BFD Session Name	LBFS_10_06_icor-to-1.1.1.1_E
BFD Session Error Code	0
AIS Monitoring	Disabled

The following example creates a static co-routed transit TP-tunnel.

```
gmpls tp-tunnel create static-transit-corout co-10.10.10.10-to-1.1.1.1 dest-  
ip 1.1.1.1 src-ip 10.10.10.10 next-hop-ip 10.11.12.1 prev-hop-ip 5.1.1.10  
forward-out-label 2001 forward-in-label 2003 reverse-out-label 2002 reverse-  
in-label 2000
```

The following sample output confirms the configuration.

```
gmpls tp-tunnel show static-transit-corout co-10.10.10.10-to-1.1.1.1
```

-----GMPLS TRANSIT TP-TUNNEL DETAILS-----	
Parameter	Value
Tunnel Name	co-10.10.10.10-to-1.1.1.1
Tunnel Index	1
Tunnel Type	Static
Direction	Bidir
Nodal Role	Transit
Destination IP Address	1.1.1.1
Source IP Address	10.10.10.10
Next-Hop IP Address	10.11.12.1
Prev-Hop IP Address	5.1.1.10
Forward In-Label	2003
Forward Out-Label	2001
Reverse In-Label	2000
Reverse Out-Label	2002
Admin State	Enabled
Oper State	Enabled
Forward CoS Profile Name	DefaultTunlCoSProfile
Forward CoS Profile Index	1
TTL Policy	decrement
Fixed TTL	255
AIS Monitoring State	Disabled

The following example creates a static co-routed egress TP-tunnel.

```
gmpls tp-tunnel create static-egress-corout ecor-frn-2.2.2.2 src-ip 2.2.2.2  
prev-hop-ip 5.1.1.1 forward-in-label 2040 reverse-out-label 2041
```

The following sample output confirms the configuration.


```
gmpls tp-tunnel show static-egress-corout egr-cor-10.10.10.10
```

-----GMPLS STATIC EGRESS TP-TUNNEL DETAILS-----	
Parameter	Value
Tunnel Name	egr-cor-10.10.10.10
Tunnel Index	2
Tunnel Type	Static
Direction	Bidir
Nodal Role	Egress
Destination IP Address	1.1.1.1
Source IP Address	10.10.10.10
Prev-Hop IP Address	10.11.12.2
Admin State	Enabled
Oper State	Enabled
Forward In-Label	2001
Reverse Out-Label	2000
Reverse Tunnel Group Index	32774
Reverse Protection Role	Primary
Reverse Protection State	Active
Reverse Backup Tunnel Name	None Present
Reverse Tunnel Reversion	On
Reverse Tunl Reversion Time	30
Reverse CoS-Profile Name	DefaultTunlCoSProfile
Reverse CoS-Profile Index	1
Reverse TTL Policy	fixed
Reverse Fixed TTL	255
BFD Monitoring	Disabled
BFD Profile ID	2
BFD Profile Name	Active-LSP
BFD Session ID	2
BFD Session Name	LBFS_12_02_egr-cor-10.10.10_E
BFD Session Error Code	0
AIS Monitoring	Disabled

Example of corouted tunnels with backup

This example shows the creation and verification of co-routed tunnels with backup tunnels as illustrated in [Figure 8-1](#).

Figure 8-1
Co-routed tunnels with backup



S1 primary static ingress corouted tunnel configuration:

```
S1*> gmpls tp-tunnel create static-ingress-corout tunnel_to_S4 dest-ip
4.4.4.4 next-hop-ip 10.1.3.2 forward-out-label 1014 reverse-in-label 1041
bfd-monitor enable ais-monitor enable
```

S1 backup static ingress corouted tunnel configuration:

```
S1*> gmpls tp-tunnel create static-ingress-corout tunnel_to_S4B dest-ip
4.4.4.4 next-hop-ip 10.1.2.2 forward-out-label 2013 reverse-in-label 2031
backup-tunnel tunnel_to_S1 bfd-monitor enable
```

S1 primary and backup static ingress corouted tunnel verification:

```
S1*> gmpls tp-tunnel show
Flags : P -> Primary      B -> Backup
        A -> Active      S -> Standby
        L -> Lone Member R -> Recovery Group Member
        E -> Tunnel uses Explicit Path
```

-----GMPLS INGRESS TP-TUNNEL TABLE-----									
Type	Direction	Tunnel Index	Tunnel Name	Destination IP	Forward Out Label	Reverse In Label	Admin State	Oper State	Flags
Static	Bidir	3	tunnel_to_S4	4.4.4.4	1014	1041	ENA	ENA	P L A
Static	Bidir	4	tunnel_to_S4B	4.4.4.4	2013	2031	ENA	ENA	B R S

S3 primary static transit corouted configuration and verification:

```
S3*> gmpls tp-tunnel create static-transit-corout Tunnel_S1_S4 dest-ip
4.4.4.4 src-ip 1.1.1.1 next-hop-ip 10.3.4.2 prev-hop-ip 10.1.3.1 forward-out-
label 1014 forward-in-label 1014 reverse-out-label 1041 reverse-in-label
1041 ais-monitor enable
```

```
S3*> gmpls tp-tunnel show
```

-----GMPLS TRANSIT TP-TUNNEL TABLE-----									
Type	Direction	Tunnel Index	Tunnel Name	Admin State	Oper State	Forward In Label	Forward Out Label	Reverse In Label	Reverse Out Label
Static	Bidir	1	Tunnel_S1_S4	ENA	ENA	1014	1014	1041	1041

S2 backup static transit corouted tunnel configuration:

```
S2*> gmpls tp-tunnel create static-transit-corout Tunnel_S1_S4B dest-ip
4.4.4.4 src-ip 1.1.1.1 next-hop-ip 10.2.4.2 prev-hop-ip 10.1.2.1 forward-out-
label 2013 forward-in-label 2013 reverse-out-label 2031 reverse-in-label
2031 ais-monitor enable
```

```
S3*> gmpls tp-tunnel show
```

-----GMPLS TRANSIT TP-TUNNEL TABLE-----									
Type	Direction	Tunnel Index	Tunnel Name	Admin State	Oper State	Forward In Label	Forward Out Label	Reverse In Label	Reverse Out Label
Static	Bidir	1	Tunnel_S1_S4B	ENA	ENA	2013	2013	2031	2031

S4 primary static egress corouted tunnel configuration

```
S4*> gmpls tp-tunnel create static-egress-corout tunnel_from_S1 src-ip
1.1.1.1 prev-hop-ip 10.3.4.1 forward-in-label 1014 reverse-out-label 1041
bfd-monitor enable ais-monitor enable
```

S4 backup static egress corouted tunnel configuration:

```
S4*> gmpls tp-tunnel create static-egress-corout from-tunnelS1B src-ip
1.1.1.1 prev-hop-ip 10.2.4.1 forward-in-label 2013 reverse-out-label 2031
backup-tunnel tunnel_from_S1 bfd-monitor enable
```

S4 primary static egress corouted tunnel verification:

```
S4*> gmpls tp-tunnel show
```

-----GMPLS STATIC EGRESS TP-TUNNEL TABLE-----									
Type	Direction	Tunnel Index	Tunnel Name	Source IP	Forward In Label	Reverse Out Label	Admin State	Oper State	Flags
Static	Bidir	3	tunnel_from_S1	1.1.1.1	1014	1041	ENA	ENA	P L A
Static	Bidir	3	from-tunneltoS1B	1.1.1.1	2013	2031	ENA	ENA	B R S

Procedure 8-7

Displaying MPLS TE-tunnel information

Display tunnel information to confirm configuration.

You can display a list of tunnels, including:

- type
- index
- name
- destination IP address
- label
- administrative state
- operational state

Optionally, you can display details about a specific tunnel. Also, you can filter to display a list of tunnels by:

- tunnel configuration (static or dynamic)
- type (ingress or egress)
- state (up or down)
- specific static egress tunnel

Step	Action

To display all or a specific tunnel

- 1 Display all or a specific tunnel:

```
mpls tunnel show matching-lsp <matching-lsp> {persist
<static|dynamic>} {type <ingress|egress|transit>} [state
<up|down>] [source <IP address>] [destination <IP
address>] [next-hop <IP address>] [out-ip-intf <Interface
Supports Signals>] [in-label <NUMBER: 16-1044479>] [out-
```

```
label <NUMBER: 16-1044479>] [recovery
<protected|unprotected>] [role <primary|backup|locally-
repaired| active-backup>]]
```

where

matching-lsp <matching-lsp>	is the tunnel to be displayed.
persist <static dynamic>	filters by persistence.
type <ingress egress transit>	filters by tunnel type.
state <up down>	filters by operational state.
source <IP address>	filters by source IP address.
destination <IP address>	filters by destination IP address.
next-hop <IP address>	filters by next hop IP address.
out-ip-intf <Interface Supports Signals>	filters by outgoing IP interface.
in-label <NUMBER: 16- 1044479>	filters by inbound label.
out-label <NUMBER: 16- 1044479>	filters by outbound label.
recovery <protected unprotected>	filters by LSP recovery type.
role <primary backup locally- repaired active- backup>	filters by LSP protection role.

To display static ingress TE-tunnels

- 2 Display static ingress TE-tunnels:

```
mpls tunnel show static-ingress <static-ingress>
```

To display static egress TE-tunnels

- 3 Display static egress TE-tunnels:

```
mpls tunnel show static-egress <static-egress> src-ip <IP  
address>
```


To display static transit TE-tunnels

- 4 Display static transit TE-tunnels:
`mpls tunnel show static-transit <static-transit>`

To display traffic statistics for static transit TE-tunnels

- 5 Display traffic statistics for static transit TE-tunnels:
`mpls tunnel show static-transit <static-transit> traffic-statistics`

To display FRR bypass TE-LSPs

- 6 Display FRR bypass TE-LSPs:
`mpls tunnel show frr-bypass-lsp <frr-bypass-lsp>`

To display bidirectional ingress associated tunnels

- 7 Display all bidirectional ingress associated TE-Tunnel(s) or only the attribute matched bidirectional associated TE tunnels:
`mpls tunnel show matching-assoc [state <up | down>]
[destination <IP address>] [next-hop <IP address>] [out-
ip-intf <Interface Supports Signals>] [out-label <NUMBER:
16-1044479>] [in-label <NUMBER: 16-1044479>] [recovery
<protected|unprotected>] [role <primary|backup|locally-
repaired|active-backup>]]`

where

matching-assoc is the tunnel to be displayed.
state <up | down> filters by operational state.
destination <IP address> filters by destination IP address.
next-hop <IP address> filters by next hop IP address.
out-ip-intf <Interface Supports Signals> filters by outgoing IP interface.
out-label <NUMBER: 16-1044479> filters by out-bound label.

where

in-label filters by in-bound label.
 <NUMBER: 16-10444795>

recovery filters by recovery type.
 <protected | unprotected>

role <primary | backup | locally-repaired | active-backup> filters by protection role.

To display bidirectional ingress associated tunnels

- 8 Display all bidirectional ingress associated TE-Tunnels:
- ```
mpls tunnel show bidir-ingress-assoc <bidir-ingress-
assoc>
```

**To display static tunnel label range information**

- 9 Display static tunnel label range information:
- ```
mpls static-tunnel-label-range show
```

To display static label range information for MPLS pseudowires

- 10 Display static label range information for MPLS pseudowires:
- ```
mpls static-vc-label-range show
```

**To display the tunnel FRR profile for a specified FRR profile**

- 11 Display the tunnel FRR profile for a specified FRR profile:
- ```
mpls tunnel-frr-profile show frr-profile <frr-profile>
```

To display the tunnel CoS profile

- 12 Display the selected tunnel CoS profile:
- ```
mpls tunnel-cos-profile show cos-profile <cos-profile>
```

—end—

## Procedure 8-8

### Displaying TP tunnel information

Display tunnel information to confirm configuration.

You can display:

- GMPLS TP tunnels
- tunnels filtered by attribute
- static ingress TP co-routed tunnels
- static ingress TP unidirectional tunnels
- static egress TP co-routed tunnels
- static egress TP unidirectional tunnels
- static transit TP co-routed tunnels
- static transit TP unidirectional tunnels
- bidirectional associated TP tunnels
- bidirectional ingress TP tunnels

| Step | Action |
|------|--------|
|------|--------|

#### *To display GMPLS TP tunnels*

- 1 Display GMPLS TP tunnels:

```
gmpls tp-tunnel show
```

#### *To display tunnels filtered by attribute*

- 2 Display tunnels filtered by attribute:

```
gmpls tp-tunnel show matching-lsp <matching-lsp> {persist
<static|dynamic>} {type <ingress|egress|transit>} [path-
type <corouted | unidirectional>] [state <up|down>]
[source <IP address>] [destination <IP address>] [next-
hop <IP address>] [prev-hop <IP address>] [fwd-out-ip-
intf <signal-interface-object>] [rev-out-ip-intf
<Interface Supports Data>] [fwd-out-ip-intf <Interface
Supports Data>] [fwd-in-label <NUMBER: 16-1044479>] [fwd-
out-label <NUMBER: 16-1044479>] [rev-in-label <NUMBER:
```

```
16-1044479>] [rev-out-label <NUMBER: 16-1044479>]
[recovery <protected|unprotected>] [role
<primary|backup|locally-repaired|active-backup>]
```

where

|                                                  |                                           |
|--------------------------------------------------|-------------------------------------------|
| matching-lsp<br><matching-lsp>                   | is the tunnel to be displayed.            |
| persist <static  <br>dynamic                     | filters by persistence.                   |
| type <ingress  <br>egress   transit>]            | filters by tunnel type.                   |
| path-type<br><corouted  <br>unidirectional>]     | filters by tunnel path type.              |
| state <up   down>                                | filters by operational state.             |
| source <IP<br>address>                           | filters by source IP address.             |
| destination <IP<br>address>                      | filters by destination IP address.        |
| next-hop <IP<br>address>                         | filters by next hop IP address.           |
| prev-hop- <IP<br>address>                        | filters by previous hop IP address.       |
| fwd-out-ip-intf<br><Interface<br>Supports Data>] | filters by forward outgoing IP interface. |
| rev-out-ip-intf<br><Interface<br>Supports Data>  | filters by reverse outgoing IP interface. |
| fwd-in-label<br><NUMBER: 16-<br>1044479>         | filters by forward in-bound label.        |
| fwd-out-label<br><NUMBER: 16-<br>1044479>        | filters by forward out-bound label.       |
| rev-in-label<br><NUMBER: 16-<br>1044479>]        | filters by reverse in-bound label.        |

where

rev-out-label filters by reverse out-bound label.  
<NUMBER: 16-1044479>

recovery filters by failure recovery type.  
<protected | unprotected>

role <primary | backup | locally-repaired | active-backup> filters by protection role.

**To display static ingress TP co-routed tunnels**

- 3 Display static ingress TP co-routed tunnels:
- ```
gmpls tp-tunnel show static-ingress-corout <static-ingress-corout>
```

To display traffic statistics for static ingress TP co-routed tunnels

- 4 Display traffic statistics for static ingress TP co-routed tunnels:
- ```
gmpls tp-tunnel show static-ingress-corout <static-ingress-corout> statistics
```

**To display static ingress TP unidirectional tunnels**

- 5 Display static ingress TP unidirectional tunnels:
- ```
gmpls tp-tunnel show static-ingress-unidir <static-ingress-unidir>
```

To display static egress TP co-routed tunnels

- 6 Display static egress TP co-routed tunnels:
- ```
gmpls tp-tunnel show static-egress-corout <static-egress-corout>
```

**To display traffic statistics for static egress TP co-routed tunnels**

- 7 Display traffic statistics for static egress TP co-routed tunnels:
- ```
gmpls tp-tunnel show static-egress-corout <static-egress-corout> statistics
```

To display static egress TP unidirectional tunnels

- 8 Display static egress TP unidirectional tunnels:
- ```
gmpls tp-tunnel show static-egress-unidir <static-egress-unidir>
```

**To display static transit TP co-routed tunnels**

- 9 Display static transit TP co-routed tunnels:
- ```
gmpls tp-tunnel show static-transit-corout <static-transit-corout>
```

To display traffic statistics for static transit TP co-routed tunnels

- 10 Display traffic statistics for static transit TP co-routed tunnels:
- ```
gmpls tp-tunnel show static-transit-corout <static-transit-corout> statistics
```

**To display static transit TP unidirectional tunnels**

- 11 Display static transit TP unidirectional tunnels:
- ```
gmpls tp-tunnel show static-transit-unidir <static-transit-unidir>
```

To display bidirectional associated TP tunnels

- 12 Display bidirectional associated TP tunnels:
- ```
gmpls tp-tunnel show matching-assoc <matching-assoc>
[state <up | down>] [destination <IP address>] [next-hop
<IP address>] [out-ip-intf <Interface Supports Signals>]
[out-label <NUMBER: 16-1044479>] [in-label <NUMBER: 16-1044479>]
[recovery <protected | unprotected>] [role
<primary | backup | locally-repaired | activebackup>]
```

**To display bidirectional ingress TP tunnels**

- 13 Display bidirectional ingress TP tunnels:
- ```
gmpls tp-tunnel show static-ingress-assoc <static-ingress-assoc>
```

—end—

Examples

The following example shows a GMPLS Static Uni-dir All TP-Tunnel Filtered Display.

```
gmpls tp-tunnel show matching-lsp persist static path-type unidirectional
```

```
Flags : P -> Primary      B -> Backup
        A -> Active      S -> Standby
        L -> Lone Member R -> Recovery Group Member
        E -> Tunnel uses Explicit Path
```

-----GMPLS INGRESS TP-TUNNEL TABLE-----									
Type	Direction	Tunnel Index	Tunnel Name	Destination IP	Forward Out Label	Reverse In Label	Admin State	Oper State	Flags
Static	Unidir	4	st-ing-u-A	1.1.1.1	1003	-	ENA	ENA	P L A
Static	Unidir	5	st-ing-u-B	2.2.2.2	1051	-	ENA	ENA	P L A

-----GMPLS STATIC EGRESS TP-TUNNEL TABLE-----									
Type	Direction	Tunnel Index	Tunnel Name	Source IP	F	In Label	Out Label	State	State
Static	Unidir	1	st-egr-u-A	1.1.1.1	1002	-	-	ENA	ENA
Static	Unidir	2	st-egr-u-B	2.2.2.2	1050	-	-	ENA	ENA

No Transit TP-Tunnel entries found

The following examples show traffic statistics for static ingress and egress TP co-routed tunnels:

```
gmpls tp-tunnel show static-ingress-corout TP_ingStatic_co_t_1 statistics
```

-----GMPLS INGRESS TP-TUNNEL TRAFFIC STATISTICS-----	
Parameter	Value
Tunnel Name	TP_ingStatic_co_t_1
Forward Tunnel Index	262145
Forward Rx Packets	2986732
Forward Tx Packets	2986774
Forward Rx Bytes	280365492
Forward Tx Bytes	280369440

```
gmpls tp-tunnel show static-egress-corout TP_egrStatic_co_t_1 statistics
```

-----GMPLS EGRESS TP-TUNNEL TRAFFIC STATISTICS-----	
Parameter	Value
Tunnel Name	TP_egrStatic_co_t_1
Reverse Tunnel Index	466945
Reverse Rx Packets	29874570
Reverse Tx Packets	29875026
Reverse Rx Bytes	2807823148
Reverse Tx Bytes	2807866012

Note 1: In the traffic statistics, the bidirectional forwarding detection (BFD) frames are accounted for in the Rx statistics at the LSP, if BFD is enabled.

Note 2: Static co-routed tunnel Tx statistics shall not be accounted if there is no VC attached to the tunnel.

Note 3: Static co-routed tunnel Rx statistics shall be accounted even though the next-hop-adjacency for that tunnel is down.

Note 4: MPLS OAM and BFD packets shall not be accounted in Tx statistics for a given static co-routed tunnel.

Note 5: Clearing the statistics of a given static co-routed tunnel shall clear all the VC's statistics (Tx/Rx) attached to that particular co-routed tunnel and vice versa.

Note 6: Tx/Rx Statistics fetch shall not be allowed on the backup tunnel (Tx/Rx statistics will be displayed on Primary LSP only which will be the cumulative sum of both primary and backup LSP).

Note 7: Statistics are supported on 5160, 3942, 5142, 3928 and 3926 devices only.

Procedure 8-9

Modifying tunnel attributes

Modify tunnel attributes if you need to change the behavior of a tunnel after initial configuration.

You can modify tunnel attributes for:

- static ingress tunnel
- associated bidirectional tunnel
- static ingress co-routed tunnel
- static egress co-routed tunnel
- static ingress unidirectional tunnel
- static egress unidirectional tunnel
- static associated bidirectional tunnel

Step	Action

To modify configuration parameters for a static ingress tunnel

- 1 Modify configuration parameters for a static ingress tunnel:

```
mpls tunnel set static-ingress <MplsIngressStaticTunl>
next-hop-ip <IpAddress> out-label <#16..1044479> cos-
profile <MplsTunlCosProfile> recovery-nhop-disjoint
<link|none> reversion-hold-time <#0..3600> tunnel-
reversion <on|off> [queue-group-profile <qgrp-name>]
[queue-group-instance <NUMBER: 1..65534>]
```

To modify configuration parameters for an associated bidirectional tunnel

- 2 Modify configuration parameters for an associated bidirectional tunnel:

```
mpls tunnel set bidir-ingress-assoc <MplsAssocTeTunl>
bfd-monitor <enable|disable> bfd-profile <BfdProfile>
```

To modify configuration parameters for a static ingress co-routed tunnel

- 3 Modify configuration parameters for a static ingress co-routed tunnel:

```
gmpls tp-tunnel set static-ingress-corout
<MplsIngressStaticTpTunl> [next-hop-ip <IpAddress> |
next-hop-local-interface <unnumbered-if-name>] forward-
out-label <#16..1044479> reverse-in-label <#16..1044479>
cos-profile <MplsTunlCosProfile> recovery-nhop-disjoint
<link|none> reversion-hold-time <#0..3600> tunnel-
reversion <on|off> bfd-monitor <enable|disable> bfd-
profile <BfdProfile> ais-monitor <enable|disable> ais-
profile <MplsAisProfiles>
```


To modify configuration parameters for a static egress co-routed tunnel

- 4 Modify configuration parameters for a static egress co-routed tunnel:
- ```
gmpls tp-tunnel set static-egress-corout
<MplsEgressStaticTpTunl> [prev-hop-ip <IpAddress> | prev-
hop-local-interface <unnumbered-if-name>] forward-in-
label <#16..1044479> reverse-out-label <#16..1044479>
cos-profile <MplsTunlCosProfile> reversion-hold-time
<#0..3600> tunnel-reversion <on|off> recovery-phop-
disjoint <link|none> bfd-monitor <enable|disable> bfd-
profile <BfdProfile> ais-monitor <enable|disable> ais-
profile <MplsAisProfiles>
```

**To modify configuration parameters for a static ingress unidirectional tunnel**

- 5 Modify configuration parameters for a static ingress unidirectional tunnel:
- ```
gmpls tp-tunnel set static-ingress-unidir
<MplsIngressStaticTpUniDirTunl> next-hop-ip <IpAddress>
forward-out-label <#16..1044479> cos-profile
<MplsTunlCosProfile> reversion-hold-time <#0..3600>
tunnel-reversion <on|off> recovery-nhop-disjoint
<link|none>
```

To modify configuration parameters for a static egress unidirectional tunnel

- 6 Modify configuration parameters for a static egress unidirectional tunnel:
- ```
gmpls tp-tunnel set static-egress-unidir
<MplsEgressStaticTpUniDirTunl> forward-in-label
<#16..1044479> prev-hop-ip <IpAddress>
```

**To modify configuration parameters for a static associated bidirectional tunnel**

- 7 Modify configuration parameters for a static associated bidirectional tunnel:
- ```
gmpls tp-tunnel set static-ingress-assoc
<MplsAssocTpTunl> bfd-monitor <enable|disable> bfd-
profile <BfdProfile> ais-monitor <enable|disable> ais-
profile <MplsAisProfiles>
```

—end—


Procedure 8-10

Inserting nodes in an unnumbered network

This procedure provides sample configuration for inserting nodes in an unnumbered network.

There are two nodes (A and C). They are connected through unnumbered IP interfaces and they have MPLS-TP static tunnel between them.

Figure 8-2
Two-node unnumbered network



Configuration at node A

```
interface create loopback lbk ip 1.1.1.1
vlan add vlan 10 port 10
interface create ip-interface ip10 ip unnumbered donor-interface lbk peer-ip-
address 3.3.3.3 ip-forwarding on vlan 10
gmpls create static-ingress-corout abc dest-ip 3.3.3.3 next-hop-local-
interface ip10 forward-out-label 100 reverse-in-label 101
```

Configuration at node C

```
interface create loopback lbk ip 3.3.3.3
vlan add vlan 10 port 21
interface create ip-interface ip10 ip unnumbered donor-interface lbk peer-ip-
address 1.1.1.1 ip-forwarding on vlan 10
gmpls create static-egress-corout abc src-ip 1.1.1.1 prev-hop-local-
interface ip10 forward-in-label 100 reverse-out-label 101
```

The new node B is inserted between A and C. Initially, node A and node C ports are on the same VLAN, but after node B is inserted in between them, one of them must move to a different VLAN. Assuming node A retains the older VLAN, then it does not need any more configuration changes apart from the peer IP address change. Node C must move to a different VLAN with the new port of node B, which means the user needs to delete the unnumbered IP interface and re-create it with a new VLAN. To do so, the user must delete all services configured over this unnumbered IP interface, that is, all PW and tunnels, then must re-create it with the new peer IP address.

Figure 8-3
Three-node unnumbered network



Configuration at node A

```
interface set ip-interface ip10 peer-ip-address 2.2.2.2
```

Configuration at node C

```
interface delete ip-interface ip10
vlan delete vlan 10 port 21
interface create loopback lbk ip 3.3.3.3
vlan add vlan 20 port 20
interface create ip-interface ip20 ip unnumbered donor-interface lbk peer-ip-
address 2.2.2.2 ip-forwarding on vlan 20
gmpls create static-egress-corout abc src-ip 2.2.2.2 prev-hop-local-interface
ip20 forward-in-label 100 reverse-out-label 101
```

Configuration at node B

```
interface create loopback lbk ip 2.2.2.2
vlan add vlan 10 port 10
vlan add vlan 20 port 20
interface create ip-interface ip10 ip unnumbered donor-interface lbk peer-ip-
address 1.1.1.1 ip-forwarding on vlan 10
interface create ip-interface ip20 ip unnumbered donor-interface lbk peer-ip-
address 3.3.3.3 ip-forwarding on vlan 20
gmpls create static-transit-corout abc dest-ip 3.3.3.3 src-ip 1.1.1.1 next-
hop-local-interface ip10 prev-hop-local-interface ip20 forward-in-label 100
forward-out-label 100 reverse-in-label 101 reverse-out-label 101
```

Confirmation of node A configuration

```
3930*> gmpls tp-tunnel show static-ingress-corout abc
```

-----GMPLS INGRESS TP-TUNNEL DETAILS-----	
Parameter	Value
Tunnel Name	abc
Tunnel Index	1
Tunnel Type	Static
Direction	Bidir
Nodal Role	Ingress
Destination IP Address	3.3.3.3
Source IP Address	1.1.1.1
Next-Hop IP Address	2.2.2.2
Next-Hop Local Iface	ip10
Next-Hop IF_NUM	0

Local IF_NUM (Next-Hop)	0
Admin State	Enabled
Oper State	Enabled
Forward Out-Label	100
Reverse In-Label	101
Forward Tunnel Group Index	32769
Forward Protection Role	Primary
Forward Protection State	Active
Forward Recovery Disjoint	link
Forward Backup Tunnel Name	None Present
Forward Tunnel Reversion	On
Forward Reversion Hold-Time	30
Forward CoS Profile Name	DefaultTunlCoSProfile
Forward CoS Profile Index	1
TTL Policy	fixed
Fixed TTL	255
BFD Monitoring	Disabled
BFD Profile ID	2
BFD Profile Name	Active-LSP
BFD Session ID	1
BFD Session Name	LBFS_10_000100_abc
BFD Session Error Code	0
AIS Monitoring	Disabled

Confirmation of node C configuration

```
3930*> gmpls tp-tunnel show static-egress-corout abc
```

-----GMPLS STATIC EGRESS TP-TUNNEL DETAILS-----	
Parameter	Value
Tunnel Name	abc
Tunnel Index	1
Tunnel Type	Static
Direction	Bidir
Nodal Role	Egress
Destination IP Address	1.1.1.1
Source IP Address	3.3.3.3
Prev-Hop IP Address	2.2.2.2
Prev-Hop Local Iface	ip20
Prev-Hop IF_NUM	0
Local IF_NUM (Prev-Hop)	0
Admin State	Enabled
Oper State	Enabled
Forward In-Label	100
Reverse Out-Label	101
Reverse Tunnel Group Index	32770
Reverse Protection Role	Primary
Reverse Protection State	Active
Reverse Recovery Disjoint	link
Reverse Backup Tunnel Name	None Present
Reverse Tunnel Reversion	On
Reverse Tunl Reversion Time	30
Reverse CoS-Profile Name	DefaultTunlCoSProfile
Reverse CoS-Profile Index	1
Reverse TTL Policy	fixed
Reverse Fixed TTL	255
BFD Monitoring	Disabled

8-38 Static tunnel configuration

BFD Profile ID	2
BFD Profile Name	Active-LSP
BFD Session ID	2
BFD Session Name	LBFS_12_000100_abc
BFD Session Error Code	0
AIS Monitoring	Disabled

Confirmation of node B configuration

```
5160*> gmpls tp-tunnel show static-transit-corout abc
```

-----GMPLS TRANSIT TP-TUNNEL DETAILS-----	
Parameter	Value
Tunnel Name	abc
Tunnel Index	5
Tunnel Type	Static
Direction	Bidir
Nodal Role	Transit
Destination IP Address	3.3.3.3
Source IP Address	1.1.1.1
Next-Hop IP Address	3.3.3.3
Next-Hop Local Iface	ip10
Prev-Hop IP Address	1.1.1.1
Prev-Hop Local Iface	ip20
Next-Hop IF_NUM	0
Prev-Hop IF_NUM	0
Local IF_NUM (Next-Hop)	0
Local IF_NUM (Prev-Hop)	0
Forward In-Label	100
Forward Out-Label	100
Reverse In-Label	101
Reverse Out-Label	101
Admin State	Enabled
Oper State	Disabled
Forward CoS Profile Name	DefaultTunlCoSProfile
Forward CoS Profile Index	1
TTL Policy	decrement
Fixed TTL	255
AIS Monitoring State	Disabled

Note: The address resolution protocol (ARP) for the peer is resolved dynamically or statically.

—end—


Procedure 8-11

Moving MPLS tunnels from numbered to unnumbered IP interfaces

This procedure provides sample configuration for moving MPLS tunnels from numbered to unnumbered IP interfaces.

There are two nodes (A and C). They are connected through numbered IP interfaces and they have MPLS-TP static tunnel between them.

Figure 8-4
Two-node numbered network



Configuration at node A

```
interface create loopback lbk ip 1.1.1.1
vlan add vlan 10 port 10
interface create ip-interface ip10 ip 10.1.1.1/24 ip-forwarding on vlan 10
gmpls create static-ingress-corout abc dest-ip 3.3.3.3 next-hop-ip 10.1.1.2
forward-out-label 100 reverse-in-label 101
```

Configuration at node C

```
interface create loopback lbk ip 3.3.3.3
vlan add vlan 10 port 21
interface create ip-interface ip10 ip 10.1.1.2/24 ip-forwarding on vlan 10
gmpls create static-egress-corout abc src-ip 1.1.1.1 prev-hop-ip 10.1.1.1
forward-in-label 100 reverse-out-label 101
```

Now they are connected through unnumbered IP interfaces and the MPLS-TP static tunnel is modified as shown in the following sample configuration.

Configuration at node A

```
vlan add vlan 11 port 10
interface create ip-interface ip11 ip unnumbered donor-interface lbk peer-ip-
address 3.3.3.3 ip-forwarding on vlan 11
gmpls disable static-ingress-corout abc
gmpls set static-ingress-corout abc next-hop-local-interface ip11
gmpls enable static-ingress-corout abc
```

Configuration at node C

```
vlan add vlan 11 port 21
interface create ip-interface ip11 ip unnumbered donor-interface lbk peer-ip-
address 1.1.1.1 ip-forwarding on vlan 11
gmpls disable static-egress-corout abc
gmpls set static-egress-corout abc prev-hop-local-interface ip11gmpls enable
static-egress-corout abc
```

—end—

Dynamic tunnel configuration

This chapter explains how to configure IP/MPLS tunnels on the 39XX/51XX Service Delivery, Aggregation and Virtualization Switches.

Dynamic MPLS-TP tunnels are:

- Dynamic ingress unidirectional (rsvp-ingress-unidir)
- Dynamic transit unidirectional (rsvp-transit-unidir)
- Dynamic co-routed (rsvp-ingress-corout)

Dynamic IP/MPLS tunnels are:

- Dynamic ingress (rsvp-ingress)
- Dynamic transit (rsvp-transit)

Dynamic tunnel configuration requires that you plan:

- [“Resiliency and robustness” on page 9-1](#)
- [“Bandwidth management and reservation” on page 9-13](#)
- [“MPLS QoS” on page 9-15](#)
- [“Labels” on page 9-19](#)
- [“MPLS fast reroute” on page 9-19](#)
- [“Preemption” on page 9-26](#)

Resiliency and robustness

Resiliency and robustness features are:

- [“BFD linkage to IS-IS” on page 9-2](#)
- [“Auto-backup” on page 9-2](#)
- [“Make-before-break” on page 9-3](#)
- [“LSP re-optimization” on page 9-6](#)
- [“Graceful restart helper mode” on page 9-7](#)
- [“Hardware-based IP forwarding” on page 9-10](#)
- [“TE link metric” on page 9-10](#)


- “Resource affinity” on page 9-11
- “SRLGs” on page 9-12

BFD linkage to IS-IS

IP BFD is configured for fault detection and to trigger dynamic LSP protection. IP BFD can be linked to IS-IS to update the IS-IS topology to find a new LSP path around a failure.

Figure 9-1 shows BFD linkage to IS-IS.

Figure 9-1
BFD linkage to IS-IS



Auto-backup

The dynamic co-routed bidirectional tunnel supports auto-backup protection. The protecting LSP for the primary LSP is automatically created. The primary LSP is established first. The secondary LSP signaling contains the protection object.

If auto-backup goes down or fails to signal, RSVP retry continues with the new constrained shortest path first (CSPF) calculations to find an alternate path. The new backup is based on shared risk link group (SRLG) differences from the primary LSP.

CSFP is the algorithm used by the head node of a traffic engineered tunnel. CSFP prunes the traffic engineering database (TED) of all links that do not meet the constraints specified when creating the tunnel.

SRLG establishes a relationship between a set of links that shares the same fault risk. If the primary LSP goes down, traffic switches over to the backup. The RSVP retry initially retries the same explicit route (ERO) for 150 seconds. Upon unsuccessful retries, a new ERO is calculated by the CSPF algorithm to


form an LSP around the failed node or link. This new calculation uses the SRLG values of the existing backup LSP to create an SRLG-diverse LSP path.

If the auto-backup feature is enabled for the LSP, the backup LSP is automatically created based on:

- node diversity
- SRLG (if enabled)
- bandwidth availability

Figure 9-2 shows auto-backup.

Figure 9-2
Auto-backup



Make-before-break

Make-before-break (MBB) creates an alternate for the existing LSP without disrupting the service traffic. The MBB uses the same egress LSR ID, tunnel ID and extended tunnel ID (which usually has the sender's LSR ID) of the original LSP. MBB is a service that handles LSP resize failure on an existing LSP path and LSP re-optimization.

Note: For multiple parallel operations, the CSPF database does not update instantly and the re-optimization may fail temporarily. The database corrects itself in the next re-optimization interval.

MBB uses the different LSP IDs to differentiate between the two LSPs. All other tunnel attributes are inherited with user-supplied modifications. A new CSPF calculation creates a new ERO. MBB is triggered by applications such as auto-size and LSP reoptimization. If MBB is triggered, it is triggered for both the primary and auto backup tunnel.

In co-routed bidirectional tunnels, MBB creates a forward and reverse LSP. It also means that MBB must be triggered at the initiating PE. Once a new LSP has been established, the service traffic is transferred and the original LSP is torn down. The new LSP is created with modified attributes, but without or minimal service disruption.

Note: CAC and MBB will not function when the device is configured at maximum scale.

MBB and auto-backup

MBB can be initiated simultaneously on the primary and the backup LSP. Each proceeds independently and can fail independently. MBB is first performed on the primary LSP. Once MBB LSP is operational, auto-backup is set on the primary MBB LSP. This yields SRLG-diverse auto-backup LSP to the new primary MBB LSP. The original primary and backup LSP is torn down.

Until the primary MBB is successful, a new auto-backup is not established. Until the auto-backup of the new primary MBB LSP succeeds, the service traffic is maintained on the original primary and backup LSP protection group. When the new MBB LSP protection group is ready, the service traffic is moved to the primary MBB LSP and the old LSP is torn down.

LSP switch-over coordination

As MBB creates a new LSP path, the PE needs to coordinate the switchover with a minimal disruption to service traffic. LSP protection groups are supported by creating two PW instances. One instance has the outer label of the primary tunnel, and the other has the standby tunnel. Only one PW instance is active at a time and during LSP switchover, the other PW instance is activated. This activation requires coordination between the control and data plane and is done individually in sequence. The actual switchover of a given PW is immediate resulting in little service disruption.

The completion of the switchover and the tear down of the original LSPs without traffic loss is a two-step process. The first step coordinates between the two PEs. The second step swaps out older LSPs with newer LSPs created by the MBB. The RSVP notify message is extended for switchover coordination.

MBB failure

The MBB, the primary LSP or the backup LSP can fail. [Table 9-1](#) describes MBB failure examples and describes the various stages of completion of MBB.

Table 9-1**MBB failure examples**

MBB failure example	MBB stage of completion
Primary LSP fails and service traffic switches over to backup LSP	MBB continues.
Primary LSP fails, the failure is repaired, the primary LSP came back up and wait-to-restore (WTR) expired on backup	MBB continues regardless of how service traffic moves around.
Backup LSP fails and RSVP retry attempts continue with new ERO	MBB continues.
Primary and backup LSPs fail	MBB continues while existing LSPs continue to recover. LSP switchover needs to be coordinated even though there is no active traffic on any of the original LSPs.

MBB does not provide alternative primary or backup LSPs. It is successful when both the alternate primary and backup LSP have been successfully established. For intermittent failure, MBB continues to retry with newer EROs until successful. It is up to the calling entity to abort.

MBB LSPs are temporary and until the service traffic is moved over to the MBB LSPs and the original LSPs are deleted, the original LSPs are the only actionable entities.

Considerations

While MBB is in progress, you cannot delete or modify the configuration of the original LSP as it impacts the ongoing MBB activities. It is up to the calling entity to determine when to stop retries. The following considerations apply:

- If MBB is not successful in one minute it is aborted to postpone re-attempt for the next timer expiry.
- Auto-size failure is mainly related to upsizing failure with an existing ERO and now MBB failure with a new ERO. When upsizing is triggered during the PW service provisioning, the MBB failure translates to the new service turn up failure. The auto-size failure indicates that the network resources are unavailable due to lack of capacity or stringent TE attributes. Operator intervention is necessary and should be terminated after two minutes.

When MBB LSPs have been established, active traffic is placed on the MBB primary LSP. It is never placed on the MBB standby LSP regardless of where the active service was present before the switchover.

Note: Auto-size (CAC and utilization based) and MBB will not function when the device is configured at maximum scale.

LSP re-optimization

LSP re-optimization attempts to find an improved LSP path, based on the latest network configuration. LSP re-optimization can be manually triggered or can run at user set intervals. LSP re-optimization uses MBB, which sets up a new LSP path before tearing down the original path to minimize traffic impact.

A dynamic LSP is created when configured by the user. A dynamic LSP uses the traffic engineering requirements of the LSP to find the shortest path from source to destination. It is possible that at the time the LSP is configured, the state of resource usage in the MPLS network does not allow the most optimized path. But at this point the path is established and it remains so until the LSP is either deleted or failed and retried.

Note: LSP re-optimization is disabled by default.


Network resource usage constantly changes due to additions and deletions of LSPs and services, links and nodes. Periodic rechecks are recommended to determine if an optimal LSP path is a possibility.

LSP re-optimization is done for each LSP at each node. This activity is coordinated at the network scale with staggered triggers. Minimal service disruption is possible due to switchover. This also needs to be considered to select candidate LSPs and the time of the day for triggering.

Note: LSP re-optimization and MBB will not function when the device is configured at maximum scale.

Figure 9-3 shows LSP optimization.

Figure 9-3
LSP optimization



Graceful restart helper mode

For control plane high availability, graceful restart for MPLS on 39XX/51XX devices is supported in helper mode only. When a device is configured with MPLS graceful restart, it assists neighboring devices that also support graceful restart to recover while in service.

Graceful restart helper mode behavior depends on the protocol enabled on the device:

- “IS-IS” on page 9-7
- “RSVP-TE” on page 9-8
- “T-LDP” on page 9-9

IS-IS

Graceful restart helper mode is enabled by default for IS-IS-TE.

The device that supports graceful restart:

- continues forwarding traffic even in the event of a control plane failure (restart) due to dataplane and control plane separation
- sets the restart request (RR) bit in the IIH packets to inform its neighbors that it is undergoing graceful restart
- waits to receive IIH with RA bit set from its neighbors before performing the graceful restart

If IIH with RA bit is not received from the neighbors after the maximum restart time, IIH with RR bit set is sent again to the neighbors.


The devices that support graceful restart helper:

- set the restart acknowledge (RA) bit in the IIH packets to acknowledge the receipt of a restart TLV with the RR bit set
- maintain IS-IS adjacencies mark them as being in restart mode
- wait up to the maximum time for LSP database synchronization

If the device that is restarting does not come up within the maximum restart time, the device declares that it has failed to achieve database synchronization.

Figure 9-4 shows an IS-IS restart scenario.

Figure 9-4
IS-IS restart scenario



RSVP-TE

Graceful restart helper mode RSVP-TE is disabled by default on devices configured with RSVP-TE.

If four hello intervals are missed, the device declares that communication has been lost and starts the restart timer.

The device that supports graceful restart:

- Keeps on forwarding traffic even in the event of a control plane failure (restart) due to dataplane and control plane separation
- Restart timer advertises how long a neighbor should wait to receive a Hello from the restarting router before it declares it failed.

The devices that support graceful restart helper mode:

- receive restart timer information from the graceful-restart capable device
- continue forwarding LSP traffic to the graceful-restart capable device while it restarts
- have a restart timer set to 0

After the graceful-restart capable device completes its restart:

- the device sends an RSVP hello message with destination ID=0
- the device receives an RSVP hello message with correct source ID from graceful restart helper devices
- RSVP hello messages are updated with the correct destination ID
- the recovery timer specifies how long the device holds the stale label-FEC bindings after an RSVP session has been reestablished

Graceful restart helper mode devices:

- receive hello PDUs from the restarted graceful restart device with destination 0 which indicates that the graceful restart device had a CP failure and recovered
- send back hello messages with their own source ID

T-LDP

T-LDP graceful restart helper mode is disabled by default.

When an LDP session goes down:

- Reconnect and recovery timers start
- The device retains the label-FEC bindings received via that session, but marks them as "stale".
- Stale bindings are deleted if lesser of reconnect time/max-peer-reconnect is expired

Graceful restart helper mode device:

- Receive reconnect and recovery timers from neighbors
- Reconnect time & Recovery time are 0 and cannot be set.
- Max-peer-reconnect and max-peer-recovery local timers are user configurable.

Device supporting Graceful restart:


- Sends in TLV reconnect and recovery timer values

- Reconnect Time is the time that the sender of the TLV would like the receiver of that TLV to wait after the receiver detects the failure of LDP session.
- Recovery Time carries the time the device is willing to retain its MPLS forwarding state that it preserved across the restart.
- Keeps on forwarding PW traffic during LDP session restart.

Figure 9-5 shows the T-LDP restart scenario. When the LDP session comes back up:

- if the graceful restart device was not able to preserve its forwarding state, all stale entries are deleted
- if the graceful restart device was able to preserve its forwarding state, stale entries are kept as the lesser of the Recovery Time and Max-peer-recovery

Figure 9-5
T-LDP restart scenario



Hardware-based IP forwarding

Hardware-based IP forwarding on 39XX/51XX devices improves IP packet forwarding performance and device scalability. Hardware-based IP forwarding is performed on transit traffic.

Hardware-based IP forwarding is always enabled on MPLS-capable devices. There is no CLI configuration related to hardware-based IP forwarding.

TE link metric

The TE link metric is a 32-bit value representing the cost of the link, which corresponds to link speed, for example:

- 100MB is 1000

- 1G is 100
- 10G is 10
- 100G is 1

The CSPF calculates the shortest path using the candidate links that have the smallest link metric values. The shortest path represents the least cost TE path.

By default, the 39XX/51XX switch assigns the TE link metric to an IP interface as the value 1 which represents a hop. The value 1 results in a constrained shortest path as least hop count. However, the least hop path may not be the shortest path. For example, a one-hop shortest path with a 100MB link represents a faster route than a two- or three-hop path with 10 or 100G links.


Resource affinity

Resource affinities allow the network operator to influence the link selection for a given LSP. A resource affinity is a 32 bit mask manually configured on a TE interface, that is, a TE interface can be assigned up to 32 different colors.

The network operator can establish resource affinity groupings based on any criteria that fit the network plan, for example, geography, administrative, or hardware. The head-end calculates a path that goes through or avoids links depending on the specified constraint.

Figure 9-6 shows an example where an LSP is constructed to avoid all red TE interfaces.

Figure 9-6
Example: LSP to far end that avoids all RED TE interfaces



SRLGs

SRLGs provide a mechanism to establish a relationship among links that share the same fault risk. For example, links that pass through the same fiber bundle share the risk that a single fiber cut will bring down all the links. By assigning the same SRLG value to all those links a backup link is created to avoid using the same link that the primary LSP is using. As a result, the end-to-end or FRR LSP protection group becomes more resilient.

The SRLG is a 32-bit value and represents the fault risk. Assigning the same value to two links makes the two links members of the shared risk group. The network operator can assign multiple SRLG values to a link. Each SRLG value represents a different risk. Risks include:


- link is on the same fiber
- link is local, for example, a VLAN on the same port
- link is on the same card and node

SRLG values are globally unique within a packet network. The optical layer uses SRLG values to protect and keep primary light paths divergent. When a packet interface is over Layer 0, 1, it is possible to expose those SRLG values to the packet layer. There is a risk of blindly inducting the optical SRLGs to the packet layer and increasing the size of SRLG list for a given packet interface.

SRLG values are distributed in the IGP advertisements and are part of link attributes in the TE database.

In [Figure 9-7](#), the primary LSP is computed and signaled regardless of the SRLG assigned to TE interfaces. Auto-backup LSP avoids using TE interfaces that have the same SRLG values used by primary LSP.

Figure 9-7
Example: SLRG



Bandwidth management and reservation

Bandwidth management and reservation features are:

- “Auto-bandwidth” on page 9-13
- “LSP bandwidth auto-resize modes” on page 9-14

Throughput is calculated by dividing the collected statistics by the size of the LSP times the duration from the last statistics collection. For example, if statistics are collected every 5 minutes on an LSP that is 200MB, then the throughput would be obtained statistics / 200MB * 300.

Auto-bandwidth

Auto-bandwidth enables auto-sizing of working LSPs based on the service load carried. The service load is determined by means of one of the following methods:

- If the LSP is only carrying the pseudowires, the sum of the bandwidth requested by each pseudowire at the time of the service configuration can be used.
- If the node can collect per LSP-based statistics, a periodic collection of actual usage of bandwidth is used.

Auto-sizing

Auto-sizing an LSP increases and decreases the LSP bandwidth to reflect the committed service load. The upsizing or downsizing requires resignaling of the RSVP PATH message with adjusted TSpec. Each node in the path needs to perform admission control, acquire/release data path resources, and notify the IGP-TE of updated bandwidth availability. The IGP-TE distributes updates from each node to all nodes in the network.

Auto-sizing LSPs lets the ingress LER admit more traffic than the network can service. It is another way to optimize network resource usage. This is more favorable than committing network resources when the LSP is configured based on intended usage that ends up idle due to actual traffic never filling the pipe.

When auto-bandwidth is configured, users need to be aware of the expected maximum bandwidth utilization, and the range of bandwidth within which the LSP is allowed to change based on the state of the service load. Auto-sizing is successful when the primary and backup LSP have auto-sized on the existing data path. There is no impact on the data path except now it operates with adjusted network resources.

User configuration dictates the handling of auto-size failure on the existing LSP path. MBB can also be used with new LSP path calculations; however, if the MBB option fails, user intervention is required.

Note: Bandwidth profile configuration is not supported on dynamic co-routed bidirectional tunnels.


LSP bandwidth auto-resize modes

LSP auto-bandwidth re-adjusts the RSVP-TE traffic specification (Tspec) bandwidth parameter if there is enough bandwidth on the line interface.

LSP auto-bandwidth can be

- manually triggered
- configured to run at user set intervals (for bandwidth decrease)
- configured to run when bandwidth has exceeded minimum and maximum set thresholds (for bandwidth increase or decrease)

Figure 9-8
Example auto-bandwidth configuration



LSP bandwidth auto-size modes are:

- [“Connection Admission Control \(CAC\)” on page 9-15](#)

- [“Utilization” on page 9-15](#)

Connection Admission Control (CAC)

Connection Admission Control (CAC) is the first line of defense used at the ingress PEs to ensure that only those flows are admitted in to the network whose service requirements are met by the network. Examples of service requirements are guarantee of bandwidth, latency, delay, and priority.

CAC maintains the record of committed flows and in the absence of available resources, it either refuses the admittance or allows over-subscription with the expectations that if congestion develops due to link failures, topology changes, or presence of higher priority traffic, uncommitted flows are preempted. CAC uses MAM bandwidth constraint algorithms to determine admission of an LSP. The pseudowire carries the service data which in turn are mapped to the established LSP. The service flows are classified at the ingress attachment circuit where policing, marking and shaping are applied before they are transported over the LSP.

Note: CAC and MBB will not function when the device is configured at maximum scale.

Utilization

Utilization mode relies on LSP statistics, which are available on 3942, 5142 and 5160 platforms.

Auto-sizing is performed based on performance monitoring of the collected statistics. Pseudowires can be oversubscribed over the LSP as the bandwidth specified for a pseudowire may not mean that all the bandwidth is in use all the time. During periods of inactivity, the LSP can be auto-sized down to accommodate other LSPs from other locations using the network resources. When the usage has increased, the LSP is dilated appropriately, maximizing the optimal usage of the network resources.

MPLS QoS

This section describes QoS on dynamic co-routed tunnels. MPLS QoS is configurable on the following platforms:

- 3928
- 3942
- 5142
- 5160

MPLS QoS features are:

- [“DiffServ-TE” on page 9-16](#)
- [“Class type” on page 9-16](#)

- [“Bandwidth constraint” on page 9-17](#)
- [“Queues and queue groups” on page 9-18](#)

Diffserv-TE-based QoS delivery configuration is done in two parts:

- configuring the per-hop-behavior on each IP interface
- requesting the per-hop-behavior from service configuration

The MPLS architecture on 39XX/51XX switches takes advantage of the existing per-hop-behavior for Layer 2 traffic to deliver QoS for the MPLS. The existing configuration is extended to support MPLS DiffServ TE based QoS delivery.

DiffServ-TE

Differentiated service in the traffic engineered MPLS network is extended by adding a class type and adding admission control using class types. There are eight class types defined. Each class type is allocated a portion of the bandwidth constraint at each link, as configured by the network operator. This available bandwidth is then advertised by IGP for each class type.

At the head-end, the network operator specifies the class type association for each LSP. The head-end processes the LSP configuration with specified TE attributes along with required bandwidth in the context of a class type and performs the CSPF using the TE database built by the IGP-TE. A set of candidate links are chosen which match the TE criteria as well as available bandwidth for a given class type. An ERO is prepared using the set of links that constitute the shortest path to the destination.

When RSVP PATH is signaled, each device performs admission control by executing admission control to determine bandwidth availability for the requested class type.

The allocated bandwidth is subtracted from the available pool for the class type and re-advertised by the IGP to keep all the nodes in the MPLS network updated with the latest bandwidth availability for each class type on each link. This in turn feeds back into CSPF for new LSP path calculation at the head-end.

If available bandwidth for a given class type becomes low for a given class type on a link, that link becomes unsuitable for newer LSPs for that class type and CSPF would choose a different set of links.


Class type

MPLS class type profile configuration is the amalgamation of the associated schedule and queue profiles and FCOS and RCOS profiles. Each MPLS class type profile is custom-defined and represents the DiffServ domain where all

the links bear the identical configuration to provide consistent per-hop behavior for each LSP. The MPLS class type profile is then associated with the MPLS-enabled IP interfaces.

Figure 9-9 shows a class type configuration example.

Figure 9-9
Class type configuration example



Bandwidth constraint


Bandwidth constraints is provided according to the Maximum Allocation Model (MAM). MAM characteristics are:

- the sum of the class type capacities must be less than or equal to the port capacity
- maximum reserved bandwidth per class type is calculated based on the class type bandwidth constraints and the local overbooking factor
- LSP resizing capabilities can be leveraged to minimize frequency of resize requests as services are added to LSPs. This applies for both CAC and utilization modes.

- class type capacities can be set to ensure total CAC bandwidth never exceeds a defined percentage of a port and to ensure bandwidth in each class type cannot be reserved by other class types
- no sharing of reserved bandwidth is allowed between class types

Figure 9-10 shows an example MAM configuration.

Figure 9-10
MAM example




Queues and queue groups

The network operator can divide bandwidth for each queue with a queue group. The port scheduler is configured with Strict or WFQ as required for each class type. Traffic treatment within each class type is based on queue configuration within the class type queue group.

For more information about traffic services configuration, refer to *39XX/51XX Service Delivery, Aggregation and Virtualization Switches Advanced Ethernet Configuration*.

Figure 9-11 shows a class type with queue groups defined.

Figure 9-11
Class type design with queue groups



Labels

Dynamic tunnels and virtual circuits use label exchange following Martini encapsulation without Control Word as defined in RFC4447. Martini encapsulation consists of adding an Ethernet transport header to the beginning of an incoming packet that includes the following:

- Destination MAC address (DA) of the PE is contained in the Ethernet transport header when directly connected to a device.
- Source MAC address of the Multi-tenant Unit (MTU) or PE and the provider VLAN.
- MPLS Ethernet-type and both the tunnel and virtual circuit MPLS labels.

MPLS fast reroute

MPLS fast reroute provides local repair for a dynamic unidirectional (RSVP-TE signaled) LSP if there is a failure in the active path of the LSP. Traffic is then redirected onto an alternate backup path with minimal traffic loss. During signaling of the LSP, each LSR in the LSP path also creates the alternative (backup) path to protect link to next-hop or protect next-hop itself.

Fast reroute is supported on the 3926m, 3928, 3942, 5142, and 5160 devices. Fast reroute works only with PHP enabled, meaning implicit null is enabled on the interface.

There are multiple options in fast reroute signaling. The options that are selected or available at a specific node are signaled to the head node. [Table 9-2](#) describes the key terms for fast reroute.

Table 9-2
Fast reroute signaling

Key term	Description
Protected LSP	The protected LSP is the primary LSP. It is signaled with fast reroute protection capability. This LSP is protected for any local failures in its path.
Point of local repair (PLR)	PLR supports manual and dynamic configuration of the facility bypass. PLR supports the auto creation of the facility bypass associated with a downstream link. PLR creates/prefers node bypass from PLR to NNHOP. PLR creates backup paths and shifts the traffic on backup paths if a failure occurs. By default, all nodes in the path of protected LSPs act as the PLR except the last node.

Table 9-2
Fast reroute signaling

Key term	Description
Merge point	This is the node where the backup path meets the protected LSP. The merge point can be next hop or next-next-hop. By default, all nodes in the protected LSP path act as the merge point except the head node.
Backup LSP	These are the backup path/LSP created at each hop in the protected LSP's route. The backup LSPs start from the PLR and finish at the merge point where it merges with the protected LSP. Each protected LSP at the PLR has a maximum of one backup LSP.
Facility bypass	This tunnel carries multiple backup LSPs passing through it. It acts as a carrier tunnel to multiple backup LSPs originating from the PLR and stopping at the merge point. The bypass tunnel is an independent tunnel used only for the purpose of carrying various backup LSPs. The facility bypass originates at the PLR and terminates at the merge point. The facility bypass tunnel can be created on the PLR by the user or be automatically generated by the PLR during the creation of the backup paths. The PLR can auto-generate the bypass tunnel when it cannot find a suitable bypass tunnel for the purposes of creating a backup LSP. The properties of the auto-generated bypass tunnel are based on the bypass profile configured on the PLR. The bypass profile is only used to determine the properties with which the bypass tunnel is created. The auto-generated bypass tunnel is disabled by default. The bypass tunnel is auto-deleted if there is no user of the bypass tunnel according to the auto facility bypass hold time. The user can explicitly delete the auto-generated bypass tunnel, but cannot disable it. Whenever an auto-generated bypass tunnel is deleted, whether manually or automatically, an SNMP trap is generated. The facility bypass tunnel has the following properties regardless of how it was created: • Dynamic unidirectional RSVP-TE signaled tunnel • Always non-protected • Re-optimization is disabled • Auto-size is disabled • Cannot be used to map services, for example, PWE or BGP cannot be mapped to the facility bypass tunnel

Facility backup method

The RSVP-TE fast reroute procedure supports the facility backup method to provide protection to a protection LSP. This method creates a single LSP that backs up a set of LSPs. This tunnel is the bypass tunnel. The bypass tunnel originates from the PLR and terminates at the merge point. When a protected LSP is being established, the PLR creates a backup path by using the bypass tunnel. The PLR creates a backup cross connect which redirects the protected LSP traffic onto the bypass tunnel which meets the protected tunnel at the merge point. The backup label is the label expected by the merge point of the protected LSP and is pushed inside the bypass tunnel label. At the merge point, when traffic arrives over a bypass tunnel, the outer label is the bypass tunnel and is popped out. The inner label is the protected LSP label, and the merge point processes this protected LSP label as if it was received on the protected link.


The facility backup method requires the use of global label space and the record route RSVP-TE feature.

When a failure occurs and the PLR detects it, the protected LSP traffic is redirected onto the backup path. The PLR starts signaling the path messages using the backup path. The signaled packets go as data packets inside the bypass tunnel.

Link and node protection


Fast reroute protection provides the option to protect against the failure of link to next hop (link protection) or against failure of next hop itself (node protection). Link protection is achieved by using a bypass tunnel between PLR and merge point. [Figure 9-12](#) shows link protection.

Figure 9-12
Link protection



[Figure 9-13](#) shows node protection.

Figure 9-13
Node level protection



During the signaling of fast reroute, PLR determines the type of protection desired by the protected LSP and then determines whether link or node protection is needed. If node protection cannot be provided, link protection is provided to the protected LSP, and the head node is updated. The system uses node protection by default.

Bandwidth protection

Bandwidth protection is used by protected LSPs to request guaranteed bandwidth in the backup path selected by any PLR. When this option is set, the PLR selects a bypass tunnel which has the bandwidth available as requested by the protected LSP. The PLR indicates the availability of the guaranteed bandwidth in the backup path by signaling the bandwidth protection flags in record route object (RRO) flags. If the PLR cannot find any bypass tunnel which has available bandwidth, the PLR still provides protection, but indicates to the head node that the bandwidth guaranteed is not provided by RRO flags.

Link affinity

IP interfaces are assigned attributes which are referred to as colors of interface. The interior gateway protocol (IGP) distributes the IGP state information of these interfaces to the attributes of each link. This information consists of bitmasks specifying multiple colors of the interfaces. RSVP can request Constrained Shortest Path First (CSPF) to select those IP interfaces for an LSP path which have the specified colors. The head LER can specify to PLR to select the backup path specific to a color.

Setup and hold priority

The setup priority for an LSP indicates the priority of the LSP that is being set up against the other existing LSP session. The holding priority holds the resources for the LSP session at that priority. The setup and hold priority determines which existing LSPs are preempted, as well as releasing held resources in favor of incoming, higher priority LSP sessions. The higher priority has a lower value: zero is highest priority, 7 is lowest in priority. The head LER specifies to the PLR the setup and hold priority for the creation of the backup path. This information is sent to the PLR as part of the FAST_REROUTE object. In the backup path creation PLR, select a bypass tunnel of equal or higher priority than specified in the FAST_REROUTE object.

Fast reroute protection switching

Fast reroute protection switching occurs when the PLR detects that the protection LSP's interface to NHOP has failed or the NHOP itself has failed. IP BFD monitors the interface failure between nodes. On a PLR when IP BFD detects that an interface has failed, it sends notification about the failure to other components of the node. Fast reroute takes all protected LSPs which have this failed interface in their path, and shifts their traffic on the backup path if it exists. Local protection switching is achieved with minimum traffic drop.

Fast reroute reversion

Reversion moves the traffic from a backup path to a working protected path. Global reversion is supported in this release.

Global reversion uses the Make Before Break (MBB) method to revert the traffic on the protected LSP. The PLR at the time of failure switches the traffic from the active path to the backup path. It also signals in the explicit route object to the head LER that at this PLR backup path is being used. This is indicated by a "Local Protection in Use" flag in the explicit route object message. When the head LER determines that global reversion needs to be started for a protected LSP, then it signals an alternate protected LSP and once that is established and up, it shifts the mapped services to this new LSP, and the old protected LSP is removed. The head LER uses MBB to shift services to the new LSP with minimal traffic drop. The protected LSP's RSVP signaling is shifted using the backup path.

If global reversion is not successful, it can be retried N times at an interval of S seconds before ending. The values of N and S are configurable.

Fast reroute reselection

Reselection of the facility bypass tunnel applies only if node protection was desired. If node protection is not available and link protection is provided, it is considered as a "compromise".

There is no reselection if bandwidth protection was asked and not found.

Overall expectations from reselection include:

- If node protection is not available, the fast reroute tunnel “compromises” to link protection and node bypass is triggered.
 - A Node-AutoFB_{compromise} is signaled “forever”.
- If neither node protection nor link protection is available, both node and link facility bypasses are triggered.
 - Eventually any of the above node or link facility bypasses provide protection.
 - If no auto facility bypasses are established and/or selected, they are timed out and deleted after one minute.
- Whenever a facility bypass comes up that can provide the desired node protection to “compromised” protected tunnels, reselection occurs and such tunnels give up any current backup to create backup with node protection.
- An unused Node-AutoFB_{compromise} is deleted if its “owner” has selected some other tunnel (or it is deleted) and the use-count of this Node-AutoFB_{compromise} is zero.
- Auto-creation and selection of the link facility bypass is supported when the node facility bypass auto-creation has failed.
- Fallback to link bypass and then upgrade to node bypass is supported.

OAM support

Table 9-3 lists OAM support features and how they work in MPLS fast reroute.

Table 9-3
OAM support

Feature	Fast reroute behavior
LSP ping	LSP ping is in parity with the unprotected tunnel's behavior. Before and after switchover, LSP ping continues to work. The time for -end-to-end may be different as it depends on the latency/length of the backup.
LSP traceroute	LSP traceroute continues to work before and after switchover. For traceroute run, post switchover, the hops are different. Post switchover, output are reflected as the PLR and merge point being neighbors. The hops do not contain the fault link/node or the hops of the facility bypass tunnel.
VC ping and traceroute	There is no impact on fast reroute switchover on VC ping and traceroute. It continues to function before and after switchover. There could be difference in end-to-end time based on the latency/length of the facility bypass path.
Bidirectional forwarding	The consistency of bidirectional forwarding cannot be ensured after switchover. The current implementation does not support LSP BFD over the fast reroute tunnels.

Preemption

Preemption is used as a tool to help ensure that high priority LSPs can always be routed through relatively favorable paths.

Preemption is an important feature in an MPLS network in which traffic is treated in a differentiated manner and highly important traffic may be given special treatment as compared to lower priority traffic.

In a Multiprotocol Label Switching (MPLS) Resource Reservation Protocol Traffic Engineering enabled IP, if LSP1 contends with LSP2 for resources, it may preempt LSP2 if it has a higher setup priority (lower numerical holding priority value) than LSP2.

Hard preemption is default behavior in MPLS RSVP-TE. Hard preemption provides no mechanism to allow preempted TE LSPs to be handled in a make-before-break fashion. Instead it performs a very intrusive method that can cause traffic disruption for a potentially large number of TE LSPs. Without an alternative, users can accept this limitation, or remove preemption functionality by using only one preemption priority (hold and setup priority) for all TE LSPs.

Initially, MPLS RSVP-TE was defined with support for only one method of TE LSP preemption, which immediately tears down TE LSPs, disregarding the preempted in-transit traffic. This simple but abrupt process nearly guarantees preempted traffic will be discarded, until a TE LSP is rerouted and a new data path can be established.

Soft preemption

With hard preemption, when a TE LSP is preempted, the preempting node sends the RSVP PathErr message. On receipt of the RSVP PathErr message, the head-end LSR sends RSVP PathTear messages, that would result in immediate traffic disruption for the preempted TE LSP.

Hard preemption is default. To avoid hard preemption, the head-end must set “Soft Preemption” while creating LSP.

With soft preemption, the preempting node takes away allocated bandwidth of the LSP that is being preempted and re-assigns the bandwidth to the incoming high priority LSP. The preempted LSP continues to propagate traffic without committed bandwidth resources. It then sends RSVP PathErr with the Ciena user-specific error code (RSVP error code 33) “Reroute request Soft Preemption” for the TE LSP upstream towards the head-end LSR. The head-end processes the received PathErr to perform make-before-break (MBB) to establish LSP to a new path away from the preempting node. Thus, soft preemption prevents traffic loss.

Note: In the user-specific error code, Ciena appends the Ciena-specific error code “Reroute request Soft Preemption”.

Bandwidth allocation in DiffServ

All LSRs advertise the available bandwidth for each priority (priority 0 to 7). For example, a 10G link advertises 10G as available for all priorities. As and when incoming LSP comes in with bandwidth for a given Setup Priority, receiving node subtracts the available bandwidth in a tiered fashion. That is if


an LSP x requests the bandwidth of 4G at Setup Priority 4, receiving node will deduct 4G from the available 10G bandwidth (as used in the example) from priority 4, 5, 6 and 7. It will then advertise updated bandwidth in the IGP-TE as 10G for priorities 0-3 but 6G for priorities 4-7.

This allows higher priority LSPs to take away bandwidth from lower priority LSPs. Again, using the same example, if an incoming LSP y with Setup Priority 2 requests 8G bandwidth, the node will take away 2G bandwidth from priority 4 LSP x, send soft preemption notification to LSP x's head-end and update the available bandwidth as 10G for priority 0 and 1 and 2G for priority 2-7.

Modes of operation

The following examples demonstrate soft preemption modes of operation.

Figure 9-14
Mode of operation - Example 1




Bandwidths and costs are identical in both directions.

Two TE tunnels are configured with soft preemption:

- LSP1 – 155 MB, setup/holding priority 0, Path R0-R1-R5
- LSP2 – 155 MB, setup/holding priority 7, Path R2-R1-R4


If the link between R1-R5 fails, R1 detects failure and sends the PathErr message to all the head-end LSRs that have a TE LSP using the failed link. See R0 in [Figure 9-15 on page 9-29](#). R0 tears down the existing LSP, triggers the reroute of LSP1, and tries to re-signal LSP1 over the shortest path available satisfying TE LSP constraints, path R0-R1-R4-R5.

Figure 9-15
Mode of operation - Example 2




The Resv messages for LSP1 travel in the upstream directions, from destination to head-end. Since the R1-R4 link does not have enough resources for LSP1, LSP2 is soft preempted at R1 because it has a numerically higher priority value.

Figure 9-16
Mode of operation - Example 3



Instead of sending a PathTear message for LSP2 on preemption as done with hard preemption, resulting in immediate traffic disruption for LSP2, R1's local bandwidth accounting is zeroed and R1 sends a PathErr message to R2, with the Ciena user-specific error code for LSP2.

Figure 9-17
Mode of operation - Example 4




On reception of PathErr, R2 re-calculates a new path excluding R1 and performs a rerouting of TE LSP2 using the make-before-break (MBB) procedure.

Once MBB succeeds, the old path is torn down. LSP2 is created on path R2-R3-R5-R4.

LSP1 is created along path R0-R1-R4-R5.

Figure 9-18
Mode of operation - Example 5



Procedure on a soft-preempting LSR

When a new TE LSP is signaled that requires a set of TE LSP or LSPs to be preempted because not all TE LSPs can be accommodated on a specific interface, a node triggers a preemption action that consists of selecting a set of TE LSPs that must be preempted so as to free up some bandwidth to satisfy the newly signaled numerically lower preemption priority (high priority) TE LSP. The preempting node may preempt TE LSPs that have a numerically higher holding priority than the setup priority of the newly signaled LSP.

The preempting node's local bandwidth accounting for each soft-preempted TE LSP is zeroed and the preempting node must immediately send a PathErr with the Ciena user-specific error code for each soft-preempted TE LSP.

To guard against a situation where bandwidth under-provisioning lasts forever, a local timer—soft preemption timer is started on the preemption node. If the preempted LSP is not removed by the head-end before this timer expires, the preempting node must perform hard preemption by sending RSVP PathTear and either a ResvTear message or a PathErr to complete removal of the LSP. The default value of the local soft preemption timer is 30 seconds.

Selection of the preempted TE LSP at a preempting midpoint: when a numerically lower value priority (that is, high priority) TE LSP is signaled that requires the preemption of a one or more numerically higher value priority (that is, lower priority) LSPs, the node where preemption is to occur has to make a decision on the set of TE LSP/LSPs that are candidates for preemption. This decision is a local one.

Procedure on head-end of soft-preempted LSP

On receipt of a PathErr message with the Ciena user-specific error code, the head-end LSR performs CSPF path computation excluding the node, which has triggered soft preemption and reroutes the TE LSP using MBB. In this case, if MBB is failed or CSPF is not able to find a new path, hard preemption must be triggered and the head-end tears down the TE LSP by sending PathTear.

Preemption policy

This section describes the criteria used for preemption.

A new LSP setup request has two important parameters for preemption: bandwidth and preemption priority. The set of LSPs to be preempted can be decided based on these two parameters.

- Preempt the LSPs that have the lowest hold priority of all.
- Preempt the least number of LSPs, so the number of LSPs that need to be rerouted is lower.
- Preempt the least amount of bandwidth that satisfies the new LSP request.

The algorithm to decide which LSPs are to be preempted is local.

Preemption for facility-bypass tunnels

A facility-bypass tunnel that is used as a backup for FRR-enabled LSPs must not be preempted because it leaves protection unavailable for the protected LSPs. This policy is enforced by virtue of signaling the facility-bypass tunnel with setup and hold priority with value zero (that is, the highest priority).

The primary LSPs that are using FRR protection can be soft preempted as described above. The head-end performs make-before-break (MBB) with a request for FRR protection away from the preempting node.

It is possible to preempt the facility-bypass tunnel but is not considered for the initial release of this feature due to complications described above. The user cannot set the hold priority value other than zero.

Configuration examples

To see sample configuration for soft preemption, take a look at:

- [“Configuring the soft preemption timer on a node” on page 9-112](#)
- [“Creating an MPLS-TE unidirectional dynamic LSP with soft preemption” on page 9-113](#)
- [“Displaying advertised bandwidth for a link” on page 9-114](#)

Preemption notes

- Preemption is supported for RSVP-TE unidirectional tunnels.

- Preemption is supported for GMPLS dynamic co-routed bidirectional tunnels, with some limitations:
 - Dynamic co-routed bidirectional tunnels can preempt any RSVP-TE unidirectional tunnels.
 - A dynamic co-routed bidirectional tunnel cannot preempt another dynamic co-routed bidirectional tunnel.
 - No tunnel can preempt a dynamic co-routed bidirectional tunnel.
- Note:** Ciena recommends that you always create a dynamic co-routed bidirectional tunnel with a higher priority than an RSVP-TE unidirectional tunnel.
- Preemption is not supported for GMPLS unidirectional tunnels or for DiffServ-TE.
 - Hard preemption is the default behavior.
 - The default soft preemption timer is 30 seconds.
 - To avoid any preemption, users can create all LSPs with default setup and hold priorities or the same setup and hold priorities.
 - The default setup and hold priorities are 0 on SAOS 6.x devices. On 6500 platforms, the default setup priority is 7 and the default hold priority is 0.

Attributes for dynamic tunnel configuration

Attributes for dynamic tunnel configuration are listed in [Table 9-4](#). Note that dynamic transit and dynamic transit unidirectional tunnels can be displayed only: there are no configuration attributes associated with the rsvp-transit and rsvp-transit-unidirectional tunnel types.

Table 9-4
Attributes for dynamic tunnel configuration

Attribute	Description	Applicable tunnel type
auto-backup <on off>	Automatically creates the protecting LSP for the primary LSP. • on: for an unprotected co-routed dynamic tunnel to create protection. Note: If the primary LSP is in transit state, such as auto-sizing or make before break (MBB) is in progress, the configuration is rejected. • off: for a protected tunnel, setting auto-protect to off removes the backup LSP Note: If service traffic is active on backup, the removal of protection is rejected.	rsvp-ingress-corout
auto-size <enable disable>	Specifies whether the auto-sizing feature is enabled.	rsvp-ingress rsvp-ingress-corout
auto-size-failure <alarm mbb>	Specifies auto-size failure handling. Valid values are: • alarm: an alarm is issued to the operator who can postpone handling of auto-size causing conditions to a quiet period so that existing services are not affected. • mbb: Auto-size failure on an existing LSP path is usually handled by MBB. A new set of primary and backup LSPs are created and changed with new TSpec. The existing user data traffic is then placed on the new primary LSP. Switchover to the newer LSPs may have minor service impact due to flight packets that could either cause loss or out-of-order delivery. Note: When MBB is enabled, it may not find an alternate path for the primary or backup LSPs or both with changed TSpec. The MBB continues to retry, but it is up to auto-size failure handling logic to intervene and abort.	rsvp-ingress rsvp-ingress-corout

Attribute	Description	Applicable tunnel type
auto-size-interval <MINUTES: 5..60>	The interval time value ranges from 5 minutes to 60 minutes. When the value is 0, the interval timer is disabled and only the manual trigger is allowed. The auto-size interval verifies whether downsizing of the LSP bandwidth is required. Using a periodic interval reduces the downsizing frequency based on PW deletion and eliminates jerks due to back-to-back deleting and adding of PWs on boundary conditions to trigger a downsize followed by an upsize.	rsvp-ingress rsvp-ingress-corout
auto-size-mode <cac utilization>	Specifies the auto-size mode which sets the bandwidth required for an LSP. The auto-size mode is one of • cac: uses mpls-vc bandwidth parameter • utilization: uses LSP statistics. This option is only available on the 3928, 3942, 5142 and 5160 platforms.	rsvp-ingress rsvp-ingress-corout
auto-size-trigger	A manual trigger for auto-downsizing and upsizing the LSP eliminates waiting for the auto-size timer to figure. The manual trigger does not reset the running timer for auto-sizing.	rsvp-ingress rsvp-ingress-corout
backup-resource include-all <MPLS TE Admin Color Group>	Specifies the administrative group that contains include-all colors to constrain path selection for the backup tunnel.	rsvp-ingress-corout
backup-resource include-any <MPLS TE Admin Color Group>	Specifies the administrative group that contains include-any colors to constrain path selection for the backup tunnel.	rsvp-ingress-corout
backup-resource exclude-any <MPLS TE Admin Color Group>	Specifies the administrative group that contains exclude-all colors to constrain path selection for the backup tunnel.	rsvp-ingress-corout
backup-tunnel <MPLS ingress primary tunnel>	Specifies the list of primary tunnels.	rsvp-ingress rsvp-ingress-unidir
bfd-monitor <enable disable>	Specifies whether BFD monitoring is enabled.	rsvp-ingress-corout
bfd-profile <profile>	Specifies the BFD profile.	rsvp-ingress-corout

Attribute	Description	Applicable tunnel type
class-type <NUMBER: 0..7>	Specifies the DSTE class type (ct0-ct7).	rsvp-ingress rsvp-ingress-corout
cos-profile <CoS Profile Name>	Selects the tunnel CoS profile.	rsvp-ingress rsvp-ingress-corout rsvp-ingress-unidir
dest-ip <IP Address>}	Sets the destination IP address for the tunnel.	rsvp-ingress rsvp-ingress-corout rsvp-ingress-unidir
exclude-ip <IP address>	Sets the IP address for the FRR facility bypass.	rsvp-ingress
explicit-tunnel-path <MPLS Rsvp Path>]	Specifies the explicit path name for the tunnel.	rsvp-ingress rsvp-ingress-corout rsvp-ingress-unidir
fixed-ttl <NUMBER: 1-255>	Specifies the fixed TTL for the tunnel.	rsvp-ingress rsvp-ingress-corout rsvp-ingress-unidir
frr-profile <MPLS Tunnel FRR Profile>]	Specifies the tunnel FRR profile name.	rsvp-ingress
frr-type <string>	Specifies the FRR type.	rsvp-ingress
hold-priority <NUMBER: 0-7>]	Specifies the tunnel hold priority.	rsvp-ingress rsvp-ingress-corout rsvp-ingress-unidir

Attribute	Description	Applicable tunnel type
increment-bandwidth <Kbps: 0..1000000000>	Sets the headroom left after auto-sizing. INC may appear as unutilized bandwidth that could be better used by some other LSP. INC's main purpose is to decrease frequent upsizing. • The value for INC is determined based on the type of PWs that are mapped to the LSP and the bandwidth characteristics of those PWs. • The INC value can be modified regardless of whether auto-sizing is enabled or disabled. The changes to INC take effect on the next auto-size trigger. The INC value is used only when auto-sizing is enabled.	rsvp-ingress rsvp-ingress-corout
lsp-reopt <enable disable>	Indicates whether LSP reoptimization is enabled. LSP reoptimization is disabled by default.	rsvp-ingress rsvp-ingress-corout
lsp-reopt-interval <NUMBER: 5..60>	Specifies the LSP reoptimization interval in minutes. The default value is 0: when the interval value is 0, the timer is disabled and only the manual trigger is allowed. Note: To properly scale, configure the lsp-reopt-interval exponentially for the tunnels, instead of configuring the same value for all the tunnels.	rsvp-ingress rsvp-ingress-corout

Attribute	Description	Applicable tunnel type
max-bandwidth <Kbps: 0..1000000000>	Sets the maximum value of the LSP. Value used by the PW admission control during configuration to validate the suitability of the PW mapping to the tunnel. If adding the PW exceeds the MAX capacity of the LSP, the PW configuration is rejected. Extremely large values that exceed bandwidth capacity of all the links in the network results in unsuccessful service activation. MAX value of the tunnel is used for both the primary and backup LSP. MAX value can be modified regardless of whether auto-sizing is enabled or disabled.	rsvp-ingress rsvp-ingress-corout
min-bandwidth <Kbps: 0..1000000000>	Sets the minimum bandwidth value of the LSP. The LSP auto-sizes between MIN and MAX MIN represents the floor value for auto-size to shrink to. When the sum of the added PWs bandwidth exceeds the MIN, the LSP upsize is triggered.	rsvp-ingress rsvp-ingress-corout
path-diverse <node srlg srlg-and-node link link-and-node srlg-and-link srlg-and-link-and-node>	Sets the backup path disjointness.	rsvp-ingress-corout
protected-interface <MPLS interface>	Specifies the protected interface.	rsvp-ingress
record-route <on off>	Specifies the record route.	rsvp-ingress rsvp-ingress-corout rsvp-ingress-unidir
resource-include-all <MPLS TE Admin Color Group>	Specifies the administrative group that contains include-all colors to constrain path selection	rsvp-ingress rsvp-ingress-corout
resource-include-any <MPLS TE Admin Color Group>	Specifies the administrative group that contains include-any colors to constrain path selection.	rsvp-ingress rsvp-ingress-corout

Attribute	Description	Applicable tunnel type
resource-exclude-any <MPLS TE Admin Color Group>	Specifies the administrative group that contains exclude-all colors to constrain path selection.	rsvp-ingress rsvp-ingress-corout
reversion-hold-time <NUMBER: 0-3600>	Sets the amount of time in seconds to wait before switching from the backup to the primary tunnel after the fault on the primary is cleared (tunnel reversion). Default is 300 seconds. Applicable when tunnel reversion is turned on.	rsvp-ingress rsvp-ingress-corout rsvp-ingress-unidir
setup-priority <NUMBER: 0-7>]	Specifies the tunnel setup priority.	rsvp-ingress rsvp-ingress-corout rsvp-ingress-unidir
soft-preemption <on off>	Specifies whether an LSP can be soft-preempted so as to free up some bandwidth to satisfy the newly signaled numerically lower preemption priority (high priority) LSP. This feature is used to prevent the disruption and traffic loss caused by the default hard preemption.	rsvp-ingress rsvp-ingress-corout rsvp-ingress-unidir
sticky-lsp <on off>	Specifies whether the tunnel reroutes to a different path on failure (such as port disable, RSVP-TE disable at interface, and so on) or keeps retrying the same path. This feature is available for TE unidirectional tunnels and for TP dynamic co-routed tunnels. For TP dynamic co-routed tunnels, when sticky-lsp is enabled for primary, it is also enabled for backup tunnels. When a tunnel is created along with MBB and sticky-lsp, then MBB is not blocked. This feature is used for fault tolerance. Sticky LSP is controlled on a per LSP basis.	rsvp-ingress rsvp-ingress-corout
ttl-policy <fixed>}	Specifies the TTL policy.	rsvp-ingress rsvp-ingress-corout rsvp-ingress-unidir
tunnel-reversion <on off>	Indicates whether to switch from the backup to the primary tunnel after the fault on the primary is cleared.	rsvp-ingress rsvp-ingress-corout rsvp-ingress-unidir

Procedures

This chapter provides the following procedures:

- [“Configuring dynamic ingress TE tunnels” on page 9-42](#)
- [“Configuring dynamic ingress unidirectional TP tunnels” on page 9-43](#)
- [“Configuring a dynamic ingress IP/MPLS tunnel with FRR signaling” on page 9-45](#)
- [“Configuring resource affinities” on page 9-48](#)
- [“Configuring SRLG” on page 9-49](#)
- [“Creating a dynamic co-routed bidirectional tunnel” on page 9-50](#)
- [“Modifying a dynamic co-routed bidirectional tunnel” on page 9-55](#)
- [“Linking a dynamic co-routed bidirectional tunnel to performance monitoring” on page 9-59](#)
- [“Configuring DiffServ-TE” on page 9-61](#)
- [“Configuring BFD linkage to IS-IS” on page 9-71](#)
- [“Displaying dynamic tunnel information” on page 9-72](#)
- [“Displaying auto-sizing statistics” on page 9-77](#)
- [“Configuring RSVP-TE graceful restart helper mode” on page 9-79](#)
- [“Configuring T-LDP graceful restart helper mode” on page 9-80](#)
- [“Configuring MPLS fast reroute” on page 9-82](#)
- [“Returning MPLS fast reroute to default values” on page 9-83](#)
- [“Creating an auto facility bypass profile” on page 9-84](#)
- [“Configuring an auto facility bypass profile” on page 9-87](#)
- [“Returning an auto facility bypass profile to default values” on page 9-88](#)
- [“Deleting an auto facility bypass profile” on page 9-89](#)
- [“Creating a fast reroute profile” on page 9-90](#)
- [“Configuring a fast reroute profile” on page 9-92](#)
- [“Returning the fast reroute profile to default values” on page 9-93](#)
- [“Deleting a fast reroute profile” on page 9-94](#)
- [“Creating a fast reroute tunnel” on page 9-95](#)
- [“Creating a facility bypass tunnel” on page 9-96](#)
- [“Configuring a fast reroute tunnel” on page 9-98](#)
- [“Returning fast reroute tunnel attributes to default values” on page 9-99](#)
- [“Deleting a fast reroute tunnel” on page 9-100](#)
- [“Configuring the IP interface with TE attributes” on page 9-101](#)

- [“Returning the IP interface TE attributes to default values” on page 9-102](#)
- [“Displaying MPLS fast reroute values” on page 9-103](#)
- [“Displaying auto facility bypass profiles” on page 9-104](#)
- [“Displaying fast reroute profiles” on page 9-106](#)
- [“Displaying fast reroute protected tunnels at ingress” on page 9-108](#)
- [“Displaying fast reroute protected tunnels at transit” on page 9-110](#)
- [“Displaying IP interface information with TE attributes” on page 9-111](#)
- [“Configuring the soft preemption timer on a node” on page 9-112](#)
- [“Creating an MPLS-TE unidirectional dynamic LSP with soft preemption” on page 9-113](#)
- [“Displaying advertised bandwidth for a link” on page 9-114](#)

Procedure 9-1

Configuring dynamic ingress TE tunnels

Configure dynamic ingress TE tunnels when setting up an LER.

Step	Action
1	Create a dynamic ingress TE tunnel: <pre>mpls tunnel create rsvp-ingress <rsvp-ingress> dest-ip <ip-address> [setup-priority <0-7>] [hold-priority <0-7>] class-type <NUMBER: 0..7> resource-include-all <MPLS TE Admin Color Group> resource-include-any <MPLS TE Admin Color Group> resource-exclude-any <MPLS TE Admin Color Group> [record-route <on off>] [frr-profile <MPLS Tunnel FRR Profile>] [explicit-tunnel-path <MPLS Rsvp Path>] [cos-profile <MPLS Tunnel COS Profile>] {ttl- policy <fixed>} [fixed-ttl <NUMBER: 1-255>] [backup- tunnel <MPLS ingress primary tunnel>] [reversion-hold- time <NUMBER: 0-3600>] [tunnel-reversion <on off>] frr- type <String> min-bandwidth <Kbps: 0..1000000000> max- bandwidth <Kbps: 0..1000000000> increment-bandwidth <Kbps: 0..1000000> auto-size-interval <MINUTES: 5..60> auto-size <String> auto-size-failure <String> protected- interface <MPLS Interface> auto-size-mode <String> lsp- reopt <String> lsp-reopt-interval <NUMBER: 5..60> sticky- lsp <on off> soft-preemption <on off> exclude-ip <IP address></pre>

—end—

Procedure 9-2

Configuring dynamic ingress unidirectional TP tunnels

Configure dynamic ingress unidirectional TP tunnels.

Step	Action
1	Create a dynamic ingress unidirectional GMPLS TP-Tunnel: <pre>gmpls tp-tunnel create rsvp-ingress-unidir <rsvp-ingress-unidir> {dest-ip <IP address>} [setup-priority <NUMBER: 0-7>] [hold-priority <NUMBER: 0-7>] [record-route <on off>] [explicit-tunnel-path <MPLS Rsvp Path>] [backup-tunnel <MPLS ingress tp-unidirtunnel>] [cos-profile <MPLS Tunnel COS Profile>] {ttl-policy <fixed>} [fixed-ttl <NUMBER: 1-255>] [reversion-hold-time <NUMBER: 0-3600>] [tunnel-reversion <on off>]</pre>
2	Confirm the configuration of the dynamic unidirectional TP tunnel: <pre>gmpls tp-tunnel show rsvp-ingress-unidir <rsvp-ingress-unidir></pre>

—end—

Example

The following example configures a dynamic unidirectional ingress TP tunnel.

```
gmpls tp-tunnel create rsvp-ingress-unidir rsvp3 dest-ip 1.1.1.1
gmpls tp-tunnel show rsvp-ingress-unidir rsvp3
```

-----GMPLS INGRESS TP-TUNNEL DETAILS-----	
Parameter	Value
Tunnel Name	rsvp3
Tunnel Index	4
Tunnel Type	Dynamic
Direction	Unidir
Nodal Role	Ingress
Destination IP Address	1.1.1.1
Source IP Address	2.2.2.2
Next-Hop IP Address	6.1.1.10
Admin State	Enabled
Oper State	Enabled
Forward Out-Label	8195
LSP ID	1
Explicit Path Index	0
Explicit Path Name	
Setup Priority	0
Hold Priority	0
Record Route	On
CSPF Route Selection	On

9-44 Dynamic tunnel configuration

Bandwidth-Profile Name	
Forward Tunnel Group Index	32772
Forward Protection Role	Primary
Forward Protection State	Active
Forward Backup Tunnel Name	None Present
Forward Tunnel Reversion	On
Forward Reversion Hold-Time	30
Forward CoS Profile Name	DefaultTunlCoSProfile
Forward CoS Profile Index	1
TTL Policy	fixed
Fixed TTL	255

+GMPLS Actual Route TABLE+	
Index	Node Ip
1	6.1.1.10
2	42.1.1.15

The following example configures a dynamic unidirectional ingress TP tunnel with backup protection.

```
gmpls tp-tunnel create rsvp-ingress-unidir rsvp-itnl-1.1.1.1 dest-ip 1.1.1.1
explicit-tunnel-path path-prim-1.1.1.1
```

```
gmpls tp-tunnel create rsvp-ingress-unidir rsvp-itnl-tp-bkp dest-ip 1.1.1.1
explicit-tunnel-path path-protect-1.1.1.1 backup-tunnel rsvp-itnl-1.1.1.1
gmpls tunnel show
```

Flags : P -> Primary B -> Backup
A -> Active S -> Standby
L -> Lone Member R -> Recovery Group Member
E -> Tunnel uses Explicit Path

-----GMPLS INGRESS TP-TUNNEL TABLE-----							
Type	Tunnel Index	Tunnel Name	Destination IP	Out Label	Admin State	Oper State	Flags
Dynamic	13	rsvp-itnl-1.1.1.1	1.1.1.1	8202	ENA	ENA	P R A E
Dynamic	14	rsvp-itnl-tp-bkp	1.1.1.1	8193	ENA	ENA	B R S E

Procedure 9-3

Configuring a dynamic ingress IP/MPLS tunnel with FRR signaling

Configure a dynamic ingress tunnel with FRR signaling to provide quick failover to a bypass LSP at an intermediate LSR when a local fault is detected. The head-end router signals FRR preferences to Point-of-Local-Repair (PLR) LSRs.

Note: FRR signaling is configurable for MPLS-TE LSPs only.

Unidirectional MPLS-TE tunnels with FRR are used for

- Ethernet pseudowires
- intra and inter-metro to access end points
- intra and inter-metro to the service PE that supports the PW head-end (PW-HE)

PW-HE is used as a logical interface by the service PE for L3VPN or L2VPN VPLS service across the MPLS core network.

Table 9-5 lists attributes for FRR profiles.

Table 9-5
FRR profile attributes

Attribute	Description
FRR profile	Tunnel FRR profile name.
Logical ID	Profile index.
Setup priority	FRR tunnel setup priority.
Hold priority	FRR tunnel hold priority.
Hop limit	FRR tunnel hop-limit.
Bandwidth	Bandwidth (Kbps/s).
Bandwidth protection	Bandwidth protection of FRR tunnel.
Node protection	Node protection of FRR tunnel.
Protection method	Bandwidth protection of FRR tunnel.
Colour group include any	Tunnel colour-group-include-any.
Colour group include all	Tunnel colour-group-include-any.
Colour group exclude any	Tunnel colour-group-include-any.

Step	Action

1	Create a tunnel Fast Reroute (FRR) profile: <pre>mpls tunnel-frr-profile create frr-profile <frr-profile> [setup-priority < NUMBER: 0-7>] [hold-priority <NUMBER: 0-7>] [hop-limit <NUMBER: 0-255>] [bandwidth <NUMBER: 1000-10000000>] [bw-protection <yes no>] [node- protection <yes no>] [color-group-include-any <NUMBER: 0-31>] [color-group-include-all <NUMBER: 0-31>] [color- group-exclude-any <NUMBER: 0-31>]</pre>
2	Confirm the FRR profile: <pre>mpls tunnel-frr-profile show mpls tunnel-frr-profile show frr-profile <frr-profile></pre>
3	Configure the tunnel at the head-end LER with FRR settings: <pre>mpls tunnel create rsvp-ingress <rsvp-ingress> dest-ip <ip-address> [setup-priority <0-7>] [hold-priority <0-7>] [bandwidth-profile <MPLS Tunnel Bandwidth Profile>] record-route <on off>][frr-profile <MPLS Tunnel FRR Profile>] [explicit-tunnel-path <MPLS Rsvp Path>] [frr-</pre>

```

type <protected | facility-bypass>][cos-profile <MPLS
Tunnel COS Profile>] {ttl-policy <fixed>} [fixed-ttl
<NUMBER: 1-255>] [backup-tunnel <MPLS ingress primary
tunnel>] [reversion-hold-time <NUMBER: 0-3600>] [tunnel-
reversion <on | off>]

```

- 4 Configure facility-based detours at each LSR (the candidate PLRs).

—end—

Example

The following example creates an IP/MPLS FRR signaling profile.

```

mpls tunnel-frr-profile create frr-profile TETUN_FRR_PROFILE node-protection
yes setup-priority 5 hold-priority bw-protection yes bandwidth 5000 colour-
group-include-any 13 colour-group-exclude-any 23

```

The following example shows sample output for the mpls tunnel-frr-profile show command.

```
mpls tunnel-frr-profile show
```

-----MPLS Tunnel Fast-Reroute Profile Table-----				
Profile Name	Protection		Protection	UseCnt
	BW	Node	Method	
DefaultFrrProfile	NO	NO	facility	0
TETUN_FRR_PROFILE	YES	YES	facility	0

```
mpls tunnel-frr-profile show frr-profile TETUN_FRR_PROFILE
```

-----MPLS Tunnel Fast-Reroute Profile Details-----	
Profile Name	TETUN_FRR_PROFILE
Profile index	2
Profile use count	0
Setup Priority	5
Hold Priority	5
Hop Limit	14
Bandwidth	5000
Bandwidth Protection	YES
Node Protection	YES
Protection Method	Facility
Color Group Include Any	13
Color Group Include All	0
Color Group Exclude Any	23

Procedure 9-4

Configuring resource affinities

Configure resource affinities to create a resource color group that can be used in the configuration of dynamic co-routed bidirectional tunnels and dynamic ingress tunnels.

Step	Action
1	Create the color-mask mapping <pre>mpls traffic-eng create resource-color <color-name> index <1..32></pre>
2	Create/Define which color(s) are part of color group <pre>mpls traffic-eng create resource-color-group <color-group-name> colors <color-name-list></pre>
3	Assign the color group to a TE interface: <pre>mpls traffic-eng set ip-interface <ip-interface-name> resource-color-group <color-group-name></pre> —end—

Procedure 9-5 Configuring SRLG

Configure an SRLG to protect LSPs that share a risk.

Ciena recommends that SRLG values are consistent across all devices in the MPLS network.

Step	Action
1	Assign SRLG values to TE interfaces: <pre>mpls traffic-eng set ip-interface <ip-interface-name unnum-ip-interface-name> srlg <32-bit-val1, ..., 32-bit- valn></pre>
2	Create primary LSP with auto-backup on: <pre>gmpls tp-tunnel create rsvp-ingress-corout <tunnel-name> dest-ip <ip-addr> auto-backup on</pre> —end—

Procedure 9-6

Creating a dynamic co-routed bidirectional tunnel

Create a dynamic co-routed bidirectional tunnel when the network plan calls for congruent forward and reverse LSPs as a single LSP unit.

A dynamic co-routed bidirectional tunnel is configured at the initiator PE. There is no tunnel related configuration at the terminating PE. When the PATH message arrives at the terminating PE, the tunnel is created automatically and the reverse path is bound to forward path.

IS-IS IGP is required for dynamic co-routed MPLS.

Note the following:

- Manual backup is not supported on dynamic co-routed bidirectional tunnels.
- Manual switchover is possible for both directions at the initiator only. It is not available at the terminator as there is no tunnel configuration.
- MBB only on auto-backup is not supported.
- A backup LSP inherits the same resource affinity constraint as the primary LSP.
- If an LSP passes through the TE link using the affinity value, the affinity value cannot be changed.

Ciena recommends that color names and masks are consistent across all devices in the MPLS network.

Step	Action

1	Create a dynamic co-routed bidirectional tunnel: <pre>gmpls tp-tunnel create rsvp-ingress-corout <tunnel-name> [logical-id <Number>] dest-ip <ip-addr> [setup-priority <0..7>] [hold-priority <0..7>] [class- type<NUMBER:0....7>] [resource-include-all <color-group- name>] [resource-include-any <color-group-name>] [resource-exclude-any <color-group-name>] [record-route <on off>] [explicit-tunnel-path <ero-path>] [path- diverse< node srlg srlg-and-node>] [auto-backup <on off>] [cos-profile <CoS profile>] [ttl-policy <fixed>] [fixed-ttl <1..255>] [reversion-hold-time <0..3600>] [tunnel-reversion <on off>] [bfd-monitor <enable disable>] [bfd-profile <MPLS BFD Profile>] [min- bandwidth<kbps:0....1000000000>] [max-bandwidth</pre>

```

<kbps:0....100000000>][increment-bandwidth
<kbps:0.....10000000>][auto-size-interval <MINUTES
5..60>][auto-size <enable|disable>][auto-size-mode
<cac|utilization>] [auto-size-failure <alarm|mbb>][lsp-
reopt <enable|disable>] [lsp-reopt-
interval<MINUTES:5....60>] [sticky-lsp <on | off> soft-
preemption <on | off>]

```

where

rsvp-ingress- corout <tunnel- name>	is the name of the tunnel.
[logical-id <Number>]	is the logical ID.
dest-ip <ip-addr>	is the destination IP address.
setup-priority <0..7>	is the setup priority. The default value is 0.
hold-priority <0..7>	is the hold type. The default value is 0.
class-type <NUMBER: 0....7>	is the class type. The default value is 0.
resource-include- all <color-group- name>	is the color group defined in “Configuring resource affinities” on page 9-48.
resource-include- any <color-group- name>	is the color group defined in “Configuring resource affinities” on page 9-48.
resource- exclude-any <color-group- name>	is the color group defined in “Configuring resource affinities” on page 9-48.
record-route <on off>	Specifies the record route. The default value is on.
explicit-tunnel- path <ero-path>	Specifies the explicit path name for the tunnel.

where	
path-diverse <node srlg srlg-and-node>	Specifies how auto-backup determines a diverse path for the backup path Valid values are: • node: auto-backup selects a path which is node diverse of the primary path • srlg: auto-backup selects a path which does not share risk with the primary path • srlg-and-node: auto-backup select a path which is both SRLG and node diverse from the primary path
auto-backup <on off>	Specifies whether auto-backup is enabled for the LSP. If auto-backup is set to on, a backup LSP is automatically created based on: • node diversity • SRLG (if enabled) • bandwidth availability
cos-profile <CoS profile>	Selects the tunnel CoS profile.
ttl-policy <fixed>]	Sets the time to live policy to fixed.
fixed-ttl <1..255>	Sets the fixed TTL. The default value is 255.
reversion-hold- time <0..3600>	Specifies the reversion hold time. The default value is 30.
tunnel-reversion <on off>	Specifies whether tunnel reversion is enabled. The default value is on.
bfd-monitor <enable disable>	Monitors individual LSPs and triggers LSP protection. The default value is disable.
bfd-profile <MPLS BFD Profile>	Identifies the BFD profile.
min-bandwidth <kbps:0.....1000 000000>	Sets the minimum threshold at which the system runs LSP auto-bandwidth. The default value is 0.
max-bandwidth <kbps:0.....1000 000000>	Sets the maximum threshold at which the system runs LSP auto-bandwidth. The default value is 1000000000.
increment- bandwidth <kbps:0.....1000 0000>	Sets the headroom left after auto-sizing. The default value is 0.

where

auto-size-interval <MINUTES 5..60>	Specifies the interval at which to run LSP auto-bandwidth for bandwidth decrease.
auto-size <enable disable>	Specifies whether LSP auto-bandwidth re-adjusts the RSVP-TE Tspec bandwidth parameter if there is enough bandwidth on the line interface.
auto-size-mode <cac utilization>	Specifies the auto-size-mode, which sets the bandwidth required for an LSP. Valid values are: • cac: uses mpls-vc bandwidth parameter • utilization: uses LSP statistics. This options is only available on the 3928, 3942, 5142 and 5160 platforms. The default value is cac.
auto-size-failure <alarm mbb>	Determines how the system responds If there is not enough bandwidth on the line interface. Valid values are • alarm: the system raises an alarm • mbb: the system locates a new LSP that has enough bandwidth by means of the MBB mechanism The default value is alarm.
lsp-reopt <enable disable>	Specifies whether LSP re-optimization is enabled.
lsp-reopt-interval <MINUTES: 5....60]	Sets the LSP re-optimization interval.
sticky-lsp <on off>	Specifies whether the tunnel reroutes to a different path on failure (such as port disable, RSVP-TE disable at interface, and so on) or keeps retrying the same path.
soft-preemption <on off>	Specifies whether an LSP can be soft-preempted so as to free up some bandwidth to satisfy the newly signaled numerically lower preemption priority (high priority) LSP. This feature is used to prevent the disruption and traffic loss caused by the default hard preemption.

—end—

Example

The following example configures a dynamic co-routed bidirectional tunnel. The tunnel is configured at the initiator PE only; the associated pseudowires are configured at the initiator PE and terminating PE.

```
gmpls tp-tunnel create rsvp-ingress-corout BetweenBostonAndSJ dest-ip
10.10.10.10
Mpls l2-vpn create dynamic-vc toSJ pw-id 10 peer 10.10.10.10 tp-tunnel-
ingress-corout BetweenBostonAndSJ pw-cword off status-tlv on
```

9-54 Dynamic tunnel configuration

```
Mpls l2-vpn create dynamic-vc toBoston pw-id 10 peer 20.20.20.20 tp-tunnel-  
egress-corout-dynamic BetweenBostonAndSJ pw-cword off status-tlv on
```

The following example creates a dynamic co-routed bidirectional tunnel with LSP BFD monitoring and auto-backup LSP:

```
gmpls tp-tunnel create rsvp-ingress-corout BetweenBostonAndSJ dest-ip  
10.10.10.10 bfd-monitor enable bfd-profile BFDBostonSJ auto-backup on
```

Procedure 9-7

Modifying a dynamic co-routed bidirectional tunnel

Modify a dynamic co-routed bidirectional tunnel as required by the network plan.

Step	Action

1	Modify a dynamic co-routed bidirectional tunnel:

```
gmpls tp-tunnel set rsvp-ingress-corout <tunnel-name>
[setup-priority <0..7>] [hold-priority <0..7>] [class-
type<NUMBER:0....7>] [resource-include-all< MPLS TE Admin
Color Group>] [resource-include-any< MPLS TE Admin Color
Group>] [resource-exclude-any <MPLS TE Admin Color Group>]
backup-resource-include-all <MPLS TE Admin Color Group>
backup-resource-include-any <MPLS TE Admin Color Group>
backup-resource-exclude-any <MPLS TE Admin Color Group>
[record-route <on | off>] [explicit-tunnel-path <ero-
path>] [auto-backup <on|off>] [path-diverse<
node|srlg|srlg-and-node>] [auto-backup <on | off>] [cos-
profile <CoS profile>] [ttl-policy <fixed>] [fixed-ttl
<1..255>] [reversion-hold-time <0..3600>] [tunnel-
reversion <on | off>] [bfd-monitor <enable | disable>]
[bfd-profile <MPLS BFD Profile>] [min-
bandwidth<kbps:0....1000000000>] [max-
bandwidth<kbps:0....1000000000>] [increment-
bandwidth<kbps:0.....100000000>] [auto-size-interval<MINUTES
5..60>] [auto-size<enable|disable>] [auto-size-
mode<cac|utilization>] [auto-size-failure
<alarm|mbb>] [lsp-reopt<enable|disable>] [lsp-reopt-
interval<MINUTES:5....60>] [sticky-lsp <on | off>] [soft-
preemption <on | off>]
```

where

rsvp-ingress-corout <tunnel-name>	is the name of the tunnel.
setup-priority <0..7>	is the setup priority. The default value is 0.
hold-priority <0..7>	is the hold type. The default value is 0.
class-type <NUMBER:0....7>	is the class type. The default value is 0.

where	
resource-include-all <MPLS TE Admin Color Group>	is the color group defined in “Configuring resource affinities” on page 9-48.
resource-include-any <MPLS TE Admin Color Group>	is the color group defined in “Configuring resource affinities” on page 9-48.
resource-exclude-any <MPLS TE Admin Color Group>	is the color group defined in “Configuring resource affinities” on page 9-48.
record-route <on off>	Specifies the record route. The default value is on.
explicit-tunnel-path <ero-path>	Specifies the explicit path name for the tunnel.
path-diverse< node srlg srlg-and-node>	Specifies how auto-backup determines a diverse path for the backup path Valid values are: • node: auto-backup selects a path which is node diverse of the primary path • srlg: auto-backup selects a path which does not share risk with the primary path • srlg-and-node: auto-backup select a path which is both SRLG and node diverse from the primary path
auto-backup <on off>	Specifies whether auto-backup is enabled for the LSP. If auto-backup is set to on, a backup LSP is automatically created based on: • node diversity • SRLG (if enabled) • bandwidth availability
cos-profile <CoS profile>	Selects the CoS profile
ttl-policy <fixed>]	Sets the time to live policy to fixed.
fixed-ttl <1..255>	Sets the fixed TTL. The default value is 255.
reversion-hold-time <0..3600>	Specifies the reversion hold time. The default value is 30.
tunnel-reversion <on off>	Specifies whether tunnel reversion is enabled. The default value is on.

where

bfd-monitor <enable disable>	Monitors individual LSPs and triggers LSP protection. The default value is disable.
bfd-profile <MPLS BFD Profile>	Identifies the BFD profile.
min-bandwidth <kbps:0.....1000 000000>	Sets the minimum threshold at which the system runs LSP auto-bandwidth. The default value is 0.
max-bandwidth <kbps:0.....1000 000000>	Sets the maximum threshold at which the system runs LSP auto-bandwidth. The default value is 1000000000.
increment- bandwidth <kbps:0.....1000 0000>	Sets the headroom left after auto-sizing. The default value is 0.
auto-size-interval <MINUTES 5..60>	Specifies the interval at which to run LSP auto-bandwidth for bandwidth decrease.
auto-size <enable disable>	Specifies whether LSP auto-bandwidth re-adjusts the RSVP-TE Tspec bandwidth parameter if there is enough bandwidth on the line interface.
auto-size-mode <cac utilization>	Specifies the auto-size-mode, which sets the bandwidth required for an LSP. Valid values are: • cac: uses mpls-vc bandwidth parameter • utilization: uses LSP statistics. This options is only available on the 3928, 3942, 5142 and 5160 platforms. The default value is cac.
auto-size-failure <alarm mbb>	Determines how the system responds If there is not enough bandwidth on the line interface. Valid values are • alarm: the system raises an alarm • mbb: the system locates a new LSP that has enough bandwidth by means of the MBB mechanism
lsp-reopt <enable disable>	Specifies whether LSP re-optimization is enabled.

where

lsp-reopt-interval Sets the LSP re-optimization interval.

<MINUTES:

5....60]

sticky-lsp <on | off> Specifies whether the tunnel reroutes to a different path on failure (such as port disable, RSVP-TE disable at interface, and so on) or keeps retrying the same path.

soft-preemption <on | off> Specifies whether an LSP can be soft-preempted so as to free up some bandwidth to satisfy the newly signaled numerically lower preemption priority (high priority) LSP. This feature is used to prevent the disruption and traffic loss caused by the default hard preemption.

—end—

Procedure 9-8

Linking a dynamic co-routed bidirectional tunnel to performance monitoring

Linking a dynamic co-routed bidirectional tunnel to performance monitoring provides the LSP utilization statistics collection and thresholding services for the tunnel. A performance monitoring instance for LSP thresholding is created after the tunnel is configured with the auto bandwidth feature.

For more information about performance monitoring, refer to *39XX/51XX Service Delivery, Aggregation, and Virtualization Switches Fault and Performance Management*.

Step	Action						
1	Create a dynamic co-routed bidirectional tunnel: <pre>gmpls tp-tunnel create rsvp-ingress-corout <tunnel-name> dest-ip <ip-addr> auto-backup on [min-bandwidth <kbps:0....1000000000>] [max-bandwidth <kbps:0....1000000000>] [increment-bandwidth <kbps:0.....100000000>] auto-size enable auto-size-mode utilization auto-size-interval <MINUTES 5..60> auto-size- failure <alarm mbb></pre>						
2	Create an interval profile: <pre>pm threshold create interval-profile <DURATION> sample- interval <1m 5m 10m 15m 30m 60m 24h> adjustment-interval <DURATION> adjustment-threshold-percentage <0..100> overflow-threshold-percentage <0..100> overflow-count <1..2147483647> underflow-threshold <0..2147483647> underflow-count <<1..2147483647> > alert-interval <DURATION></pre> where <table> <tr> <td>interval-profile <interval-profile-name></td><td>is the name of the interval profile.</td></tr> <tr> <td>sample-interval <1m 5m 10m 15m 30m 60m 24h></td><td>is the interval for statistics collection.</td></tr> <tr> <td>adjustment-interval <DURATION></td><td>Time interval at which determination is made whether to adjust the LSP bandwidth or not</td></tr> </table>	interval-profile <interval-profile-name>	is the name of the interval profile.	sample-interval <1m 5m 10m 15m 30m 60m 24h>	is the interval for statistics collection.	adjustment-interval <DURATION>	Time interval at which determination is made whether to adjust the LSP bandwidth or not
interval-profile <interval-profile-name>	is the name of the interval profile.						
sample-interval <1m 5m 10m 15m 30m 60m 24h>	is the interval for statistics collection.						
adjustment-interval <DURATION>	Time interval at which determination is made whether to adjust the LSP bandwidth or not						

where	
overflow-threshold-percentage <0..100>	This threshold is used to check against the DIFF(sampleBw, current-LSP-BW) to see whether bandwidth usage is rapidly increasing
overflow-count <1..2147483647>	A consecutive count the overflow-threshold was crossed. When this count is reached, LSP autosize is dispatched immediately without waiting for adjustment interval to expire
underflow-threshold <0..2147483647>	This threshold is used to check against the DIFF(SampleBw, current-LSP-BW) to see if b/w usage is rapidly decreasing
underflow-count <1..2147483647>	A consecutive count the Underflow Threshold has crossed. When this count is reached the LSP is downsized immediately to MaxAvgBw value.
alert-interval <DURATION>	is the alert interval

3 Create statistics collection:

```
pm create tp-rsvp-ingress-corout <tunnel-name> pm-
instance <pm-instance> instance-type proactive profile-
type BasicTxRx bin-duration <1m | 5m | 10m | 15m | 30m |
60m | 24h> interval-profile <internal-profile-name>
```

—end—

Example

The following example creates a dynamic co-routed bidirectional tunnel that is lined to performance monitoring.

```
gmpls tp-tunnel create rsvp-ingress-corout S3-S4_50-100+10 dest-ip 4.4.4.4
auto-backup on min-bandwidth 50000 max-bandwidth 100000 increment-bandwidth
10000 auto-size enable auto-size-mode utilization auto-size-interval 5 auto-
size-failure mbb
```

```
pm threshold create interval-profile AutoSize sample-interval 5m adjustment-
interval 10m adjustment-threshold-percentage 2 overflow-threshold-percentage
2 overflow-count 2 underflow-threshold 2 underflow-count 2 alert-interval 5m
```

```
pm create tp-rsvp-ingress-corout S3-S4_50-100+10 pm-instance AutoSizeInstance
instance-type proactive profile-type BasicTxRx bin-duration 5m interval-
profile AutoSize
```

Procedure 9-9

Configuring DiffServ-TE

Configure DiffServ-TE to meet the traffic engineering requirements of the network plan.

When configuring DiffServ-TE, note the following:

- a tunnel with a class-type greater than 0 and min-bandwidth equal to 0 will not come up
- a mix of DS and DSTE interfaces, that is, a DSTE tunnel with a class type greater than 0, may not get established due to the mix of DS and DSTE interfaces

When assigning resources, note the following:

- The mpls-queuing feature is used for directing MPLS transit and (non-stitched) L2-VPN traffic to egress queue group according to the class type profile attached to the IP interface of the tunnel. When enabled,
 - each class type-mapped transit tunnel consumes one classifier
 - each non-stitched L2-VPN over class type-mapped unprotected tunnel consumes one classifier
 - each non-stitched L2-VPN over class type-mapped protected tunnel consumes two classifiers
- The mpls-mspw-queuing feature is used for directing stitched L2-VPN traffic to egress queue group according to the class type profile attached to the IP interface of the tunnel. When enabled,
 - each stitched L2-VPN over class type-mapped unprotected tunnel consumes one classifier
 - each stitched L2-VPN over class type-mapped protected tunnel consumes two classifiers.

Table 9-6 lists groups and resources for each platform.:

Table 9-6
Groups and resources for each platform

Switch	Resource	Number of groups	Resources for each group
5142, 5160	classifiers	16	512
	meters	meter resources cannot be reserved: they are allocated from a global pool on a first-come, first-served basis when needed	global pool of 8192
	counters	counter resources cannot be reserved: they are allocated from a global pool on a first-come, first-served basis when needed	global pool of 8192
3903, 3903x, 3904, 3905, 3906mvi, 3916, 3930, 3931, 3932	classifiers	8	256
	meters	meter resources cannot be reserved: they are allocated from a global pool on a first-come, first-served basis when needed	global pool of 2048
	counters	8	256
3942, 3938	classifiers	16	256
	meters	meter resources cannot be reserved: they are allocated from a global pool on a first-come, first-served basis when needed	global pool of 4096
	counters	counter resources cannot be reserved: they are allocated from a global pool on a first-come, first-served basis when needed	global pool of 4096

Table 9-6
Groups and resources for each platform

Switch	Resource	Number of groups	Resources for each group
3926m, 3928	classifiers	8	256
	meters	meter resources cannot be reserved: they are allocated from a global pool on a first-come, first-served basis when needed	global pool of 2048
	counters	8	256

Prerequisite

Ensure that you have access to *39XX/51XX Service Delivery, Aggregation, and Virtualization Switches Advanced Ethernet Configuration*.

Step	Action

Assign resources

1	Determine whether resources are to be assigned to a single-segment pseudowire or a multi-segment pseudowire:					
	<table border="1"> <thead> <tr> <th>If resources are to be assigned to a</th><th>Then</th></tr> </thead> <tbody> <tr> <td>single-segment pseudowire</td><td> Assign classifier resources to the mpls-queuing feature: <pre>resource-manager pool set feature mpls-queuing resource classifier count <NUMBER></pre> Assign counter resources to the mpls-queuing feature: <pre>resource-manager pool set feature mpls-queuing resource counter count <NUMBER></pre> </td></tr> <tr> <td>multi-segment pseudowire</td><td> Assign classifier resources to the mpls-mspw-queuing feature: <pre>resource-manager pool set feature mpls-mspw-queuing resource classifier count <NUMBER></pre> Assign counter resources to the mpls-mspw-queuing feature: <pre>resource-manager pool set feature mpls-mspw-queuing resource counter count <NUMBER></pre> </td></tr> </tbody> </table>	If resources are to be assigned to a	Then	single-segment pseudowire	Assign classifier resources to the mpls-queuing feature: <pre>resource-manager pool set feature mpls-queuing resource classifier count <NUMBER></pre> Assign counter resources to the mpls-queuing feature: <pre>resource-manager pool set feature mpls-queuing resource counter count <NUMBER></pre>	multi-segment pseudowire
If resources are to be assigned to a	Then					
single-segment pseudowire	Assign classifier resources to the mpls-queuing feature: <pre>resource-manager pool set feature mpls-queuing resource classifier count <NUMBER></pre> Assign counter resources to the mpls-queuing feature: <pre>resource-manager pool set feature mpls-queuing resource counter count <NUMBER></pre>					
multi-segment pseudowire	Assign classifier resources to the mpls-mspw-queuing feature: <pre>resource-manager pool set feature mpls-mspw-queuing resource classifier count <NUMBER></pre> Assign counter resources to the mpls-mspw-queuing feature: <pre>resource-manager pool set feature mpls-mspw-queuing resource counter count <NUMBER></pre>					

- 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`

Configure an RCoS profile

- 5 Determine whether the default RCoS profile can be used.

If the default RCoS profile	Then
can be used	go to step 6 .
cannot be used	Perform the following procedures: • Creating a resolved CoS map • Modifying a resolved CoS map Refer to <i>39XX/51XX Service Delivery, Aggregation, and Virtualization Switches Advanced Ethernet Configuration</i> .

Configure an FCoS profile

- 6 Determine whether the default FCoS profile can be used.

If the default RCoS profile	Then
can be used	go to step 10 .
cannot be used	Perform the following procedures: • Creating a frame CoS map • Modifying a frame CoS map Refer to <i>39XX/51XX Service Delivery, Aggregation, and Virtualization Switches Advanced Ethernet Configuration</i> .

Configure scheduling

- 7 Define the scheduling on all MPLS ports. Use the procedure “Configuring the port scheduler for hierarchical egress,” in *39XX/51XX Service Delivery, Aggregation, and Virtualization Switches Advanced Ethernet Configuration*.

Create egress queue group with egress queue group profile parameters for each class type

- 8 Create an egress queue group with queue parameters for each class type. Use the procedure “Configuring the queue group profile for hierarchical egress queuing” or “Configuring the dynamic queue group profile for hierarchical egress queuing,” in *39XX/51XX Service Delivery, Aggregation, and Virtualization Switches Advanced Ethernet Configuration*.

Assign class type queue groups to port

- 9 Assign class type queue groups to the port and define queue group CIR and EIR parameters. Use the procedure “Adding a port to a customer queue group,” in *39XX/51XX Service Delivery, Aggregation, and Virtualization Switches Advanced Ethernet Configuration*.

Configure an MPLS class type profile

- 10 Create an MPLS class type profile:

```
mpls traffic-eng class-type-profile create class-profile
<class-type-profile-name>
```
- 11 Link the class type profile to the queue group:

```
mpls traffic-eng class-type-profile set class-profile
<class-type-profile-name> class-type <NUMBER: 0..7>
queue-group <queue-group-profile-name> lom < NUMBER:
1..4> ]
```

where

class-profile <class-type-profile-name> Specifies the class type profile to be modified.

class-type <NUMBER: 0..7> Specifies the class type in the format CT0 to CT7.

queue-group <queue-group-profile-name> Is the queue group to add to the class type profile.

lom < NUMBER: 1..4> Specifies the link subscription multiplier value. The default value is 1.

- 12 (Optional) Configure alarm threshold:

```
mpls traffic-eng class-type-profile set class-profile
<class-type-profile-name> class-type <NUMBER: 0..7>
alarm-threshold-percentage <NUMBER: 0..100> queue-group
<queue-group-profile-name>
```

Assign an MPLS class type profile to an interface

- 13 Assign the MPLS class type profile to the IP interface on the port on which the queue groups are applied:

```
mpls traffic-eng set ip-interface <ip-interface> class-
profile <class-profile>
```


Associate LSP to class type at the head-end node

- 14 Associate LSP to class type when creating the LSP at the head-end node:
- ```
gmpls tp-tunnel create rsvp-ingress-corout <rsvp-ingress-
corout> dest-ip <IP address> min-bandwidth
<kbps:0....1000000000> max-bandwidth
<kbps:0....1000000000> auto-size enable

gmpls tp-tunnel create rsvp-ingress-corout <rsvp-ingress-
corout> dest-ip <IP address> class-type <NUMBER: 0....7>
min-bandwidth <kbps:0....1000000000> max-bandwidth
<kbps:0....1000000000> auto-size enable
```


**Create and define pseudowire bandwidth and assign to tunnel**

- 15 Create and define pseudowire bandwidth on the end nodes and assign the pseudowire bandwidth to the tunnel:
- ```
mpls l2-vpn create dynamic-vc <dynamic-vc> pw-id <NUMBER:
1-2147483647> peer <IP address> tp-tunnel-ingr-corout
<MPLS ingress primary tp corout tunnel> pw-cword <on/off>
bandwidth <Kbps>
```
- end—

Example

In the example shown in [Figure 9-19](#), three tunnel class types with three different bandwidth constraints in conjunction with three queue group schedulers are used to limit a queue in each service. There are multiple priority queues for each service.

Figure 9-19
Example topology



Assign resources for all DIFFSERV-TE nodes:

```
resource-manager pool set feature mpls-queuing resource classifier count 1024
resource-manager pool set feature mpls-queuing resource counter count 1024
```

Configure a custom RCoS profile:

```
traffic-services cos-mapping resolved-cos-map create cos-map DS-TE-RCOS
traffic-services cos-mapping resolved-cos-map set cos-map DS-TE-RCOS mpls-tc
7 r-cos 7
```

```

traffic-services cos-mapping resolved-cos-map set cos-map DS-TE-RCOS mpls-tc
6 r-cos 6
traffic-services cos-mapping resolved-cos-map set cos-map DS-TE-RCOS mpls-tc
5 r-cos 5
traffic-services cos-mapping resolved-cos-map set cos-map DS-TE-RCOS mpls-tc
4 r-cos 4
traffic-services cos-mapping resolved-cos-map set cos-map DS-TE-RCOS mpls-tc
3 r-cos 3
traffic-services cos-mapping resolved-cos-map set cos-map DS-TE-RCOS mpls-tc
2 r-cos 2
traffic-services cos-mapping resolved-cos-map set cos-map DS-TE-RCOS mpls-tc
1 r-cos 1
traffic-services cos-mapping resolved-cos-map set cos-map DS-TE-RCOS mpls-tc
0 r-cos 0
port set port 24 resolved-cos-map DS-TE-RCOS

```

Configure a custom FCoS profile:

```

traffic-services cos-mapping frame-cos-map create cos-map DS-TE-FCOS
traffic-services cos-mapping frame-cos-map set cos-map DS-TE-FCOS r-cos 7 r-
color green mpls-tc 7
traffic-services cos-mapping frame-cos-map set cos-map DS-TE-FCOS r-cos 6 r-
color green mpls-tc 6
traffic-services cos-mapping frame-cos-map set cos-map DS-TE-FCOS r-cos 5 r-
color green mpls-tc 5
traffic-services cos-mapping frame-cos-map set cos-map DS-TE-FCOS r-cos 4 r-
color green mpls-tc 4
traffic-services cos-mapping frame-cos-map set cos-map DS-TE-FCOS r-cos 3 r-
color green mpls-tc 3
traffic-services cos-mapping frame-cos-map set cos-map DS-TE-FCOS r-cos 2 r-
color green mpls-tc 2
traffic-services cos-mapping frame-cos-map set cos-map DS-TE-FCOS r-cos 1 r-
color green mpls-tc 1
traffic-services cos-mapping frame-cos-map set cos-map DS-TE-FCOS r-cos 0 r-
color green mpls-tc 0
port set port 24 frame-cos-map DS-TE-FCOS

```

Define the scheduling on all MPLS ports:

```

traffic-services queuing egress-port set port 24 scheduler-algorithm
weighted-deficit-round-robin

```

Create Egress Queue Group with Queue Parameters for CT0:

```

traffic-services queuing dynamic-egress-queue-group-profile create profile
DS-TE-QG-CT0Pro
traffic-services queuing dynamic-egress-queue-group-profile set queue 0
profile DS-TE-QG-CT0Pro scheduler-weight 10
traffic-services queuing dynamic-egress-queue-group-profile set queue 1
profile DS-TE-QG-CT0Pro scheduler-weight 10 cir-percent 100 cbs 1024
traffic-services queuing dynamic-egress-queue-group-profile set queue 2
profile DS-TE-QG-CT0Pro scheduler-weight 10
traffic-services queuing dynamic-egress-queue-group-profile set queue 3
profile DS-TE-QG-CT0Pro scheduler-weight 10
traffic-services queuing dynamic-egress-queue-group-profile set queue 4
profile DS-TE-QG-CT0Pro scheduler-weight 10
traffic-services queuing dynamic-egress-queue-group-profile set queue 5
profile DS-TE-QG-CT0Pro scheduler-weight 10

```

```
traffic-services queuing dynamic-egress-queue-group-profile set queue 6
profile DS-TE-QG-CT0Pro scheduler-weight 0
traffic-services queuing dynamic-egress-queue-group-profile set queue 7
profile DS-TE-QG-CT0Pro scheduler-weight 0
traffic-services queuing egress-queue-group create queue-group DS-TE-QG-CT0
profile DS-TE-QG-CT0Pro
```

Create Queue Group with Queue Parameters for CT4:

```
traffic-services queuing dynamic-egress-queue-group-profile create profile
DS-TE-QG-CT4Pro
traffic-services queuing dynamic-egress-queue-group-profile set queue 0
profile DS-TE-QG-CT4Pro scheduler-weight 10
traffic-services queuing dynamic-egress-queue-group-profile set queue 1
profile DS-TE-QG-CT4Pro scheduler-weight 10 cir-percent 100 cbs 1024
traffic-services queuing dynamic-egress-queue-group-profile set queue 2
profile DS-TE-QG-CT4Pro scheduler-weight 10
traffic-services queuing dynamic-egress-queue-group-profile set queue 3
profile DS-TE-QG-CT4Pro scheduler-weight 10
traffic-services queuing dynamic-egress-queue-group-profile set queue 4
profile DS-TE-QG-CT4Pro scheduler-weight 10
traffic-services queuing dynamic-egress-queue-group-profile set queue 5
profile DS-TE-QG-CT4Pro scheduler-weight 10
traffic-services queuing dynamic-egress-queue-group-profile set queue 6
profile DS-TE-QG-CT4Pro scheduler-weight 0
traffic-services queuing dynamic-egress-queue-group-profile set queue 7
profile DS-TE-QG-CT4Pro scheduler-weight 0
traffic-services queuing egress-queue-group create queue-group DS-TE-QG-CT4
profile DS-TE-QG-CT4Pro
```

Create Queue Group with Queue Parameters for CT7:

```
traffic-services queuing dynamic-egress-queue-group-profile create profile
DS-TE-QG-CT7Pro
traffic-services queuing dynamic-egress-queue-group-profile set queue 0
profile DS-TE-QG-CT7Pro scheduler-weight 10
traffic-services queuing dynamic-egress-queue-group-profile set queue 1
profile DS-TE-QG-CT7Pro scheduler-weight 10 cir-percent 100 cbs 1024
traffic-services queuing dynamic-egress-queue-group-profile set queue 2
profile DS-TE-QG-CT7Pro scheduler-weight 10
traffic-services queuing dynamic-egress-queue-group-profile set queue 3
profile DS-TE-QG-CT7Pro scheduler-weight 10
traffic-services queuing dynamic-egress-queue-group-profile set queue 4
profile DS-TE-QG-CT7Pro scheduler-weight 10
traffic-services queuing dynamic-egress-queue-group-profile set queue 5
profile DS-TE-QG-CT7Pro scheduler-weight 0
traffic-services queuing dynamic-egress-queue-group-profile set queue 6
profile DS-TE-QG-CT7Pro scheduler-weight 0
traffic-services queuing dynamic-egress-queue-group-profile set queue 7
profile DS-TE-QG-CT7Pro scheduler-weight 0
traffic-services queuing egress-queue-group create queue-group DS-TE-QG-CT7
profile DS-TE-QG-CT7Pro
```

Assign CT Queue groups to port and define queue group CIR/EIR parameters:

```
traffic-services queuing egress-queue-group add port 24 queue-group DS-TE-QG-
CT0 priority 10 scheduler-weight 10 cir 5000000 cbs 1024
```

```
traffic-services queuing egress-queue-group add port 24 queue-group DS-TE-QG-CT4 priority 14 scheduler-weight 10 cir 4000000 cbs 1024
traffic-services queuing egress-queue-group add port 24 queue-group DS-TE-QG-CT7 priority 17 scheduler-weight 10 cir 990016 cbs 1024
```

Create MPLS CT Profile and link CTs to queue groups:

```
mpls traffic-eng class-type-profile create class-profile CT-10G
mpls traffic-eng class-type-profile set class-profile CT-10G class-type 0
queue-group DS-TE-QG-CT0
mpls traffic-eng class-type-profile set class-profile CT-10G class-type 4
queue-group DS-TE-QG-CT4 lom 2
mpls traffic-eng class-type-profile set class-profile CT-10G class-type 7
queue-group DS-TE-QG-CT7
```

Configure alarm threshold:

```
mpls traffic-eng class-type-profile set class-profile CT-10G class-type 4
alarm-threshold-percentage 75 queue-group DS-TE-QG-CT4
```

Assign MPLS CT Profile to IP Interface(s) present on the port on which the queue groups are applied:

```
mpls traffic-eng set ip-interface toS4 class-profile CT-10G
```

Associate LSP to Class Type when creating LSP on head end node.

```
gmpls tp-tunnel create rsvp-ingress-corout S3S13CT0 dest-ip 13.13.13.13 min-bandwidth 100000 max-bandwidth 5000000 auto-size enable
gmpls tp-tunnel create rsvp-ingress-corout S3S13CT7 dest-ip 13.13.13.13 class-type 7 min-bandwidth 100000 max-bandwidth 4000000 auto-size enable
gmpls tp-tunnel create rsvp-ingress-corout S3S13CT4 dest-ip 13.13.13.13 class-type 4 min-bandwidth 4000000 max-bandwidth 4000000 auto-size enable
```

Associate LSP to Class Type when creating LSP on head end node.

```
gmpls tp-tunnel create rsvp-ingress-corout S3S13CT0 dest-ip 13.13.13.13 min-bandwidth 100000 max-bandwidth 5000000 auto-size enable
gmpls tp-tunnel create rsvp-ingress-corout S3S13CT7 dest-ip 13.13.13.13 class-type 7 min-bandwidth 100000 max-bandwidth 4000000 auto-size enable
gmpls tp-tunnel create rsvp-ingress-corout S3S13CT4 dest-ip 13.13.13.13 class-type 4 min-bandwidth 4000000 max-bandwidth 4000000 auto-size enable
```

Create and define PW bandwidth on end nodes and assign it to tunnel.

S3:

```
mpls l2-vpn create dynamic-vc PW1-CT0 pw-id 100 peer 13.13.13.13 tp-tunnel-ingr-corout S3S13CT0 bandwidth 5000000
mpls l2-vpn create dynamic-vc PW1-CT4 pw-id 200 peer 13.13.13.13 tp-tunnel-ingr-corout S3S13CT4 bandwidth 4000000
mpls l2-vpn create dynamic-vc PW1-CT7 pw-id 300 peer 13.13.13.13 tp-tunnel-ingr-corout S3S13CT7 bandwidth 990000
```

S13:

9-70 Dynamic tunnel configuration

```
mpls l2-vpn create dynamic-vc PW1-CT0 pw-id 100 peer 3.3.3.3 tp-tunnel-egrs-  
corout-dynamic S3S13CT0 bandwidth 5000000  
mpls l2-vpn create dynamic-vc PW1-CT4 pw-id 200 peer 3.3.3.3 tp-tunnel-egrs-  
corout-dynamic S3S13CT4 bandwidth 4000000  
mpls l2-vpn create dynamic-vc PW1-CT7 pw-id 300 peer 3.3.3.3 tp-tunnel-egrs-  
corout-dynamic S3S13CT7 bandwidth 990000
```

Procedure 9-10

Configuring BFD linkage to IS-IS

Configure BFD linkage to IS-IS to provide protection for LSPs along the link.

Step	Action
1	Create an IS-IS instance: <pre>isis instance create isis-instance <isis-instance> level L1 area <ISIS area></pre>
2	Attach an IP interface to IS-IS instance: <pre>isis interface attach ip-interface <ip-interface> isis- instance <ISIS instance> level L1</pre>
3	Enable BFD monitoring: <pre>isis interface set ip-interface <> bfd-monitor enable</pre> —end—

Example

The following example configures BFD linkage to IS-IS for an ip-interface named mpls_3-2.

```
isis instance create isis-instance MPLS level L1 area 49.0000
isis interface attach ip-interface mpls_3-2 isis-instance MPLS level L1
isis interface set ip-interface mpls_3-2 bfd-monitor enable
```

Procedure 9-11

Displaying dynamic tunnel information

Display dynamic tunnel information to confirm configuration.

You can display a list of tunnels, including:

- type
- index
- name
- destination IP address
- label
- administrative state
- operational state

Optionally, you can display details about a specific tunnel. Also, you can filter to display a list of tunnels by:

- tunnel configuration (static or dynamic)
- type (ingress or egress)
- state (up or down)

Step	Action

To display all TE tunnels

- 1 Display all TE tunnels:
`mpls tunnel show`

To display all TP tunnels

- 2 Display all TP tunnels:
`gmpls tp-tunnel show`

To display TE tunnels filtered by attribute

- 3 Display TE tunnels filtered by attribute:
`mpls tunnel show matching-lsp <matching-lsp> {persist
<static|dynamic>} {type <ingress|egress|transit>} [state
<up|down>] [source <IP address>] [destination <IP
address>] [next-hop <IP address>] [out-ip-intf <Interface
Supports Signals>] [in-label <NUMBER: 16-1044479>] [out-
label <NUMBER: 16-1044479>] [recovery`

```
<protected|unprotected>] [role <primary|backup|locally-
repaired| active-backup>] [cos-profile <MPLS Tunnel COS
Profile>]]
```

where

matching-lsp <matching-lsp>	is the tunnel to be displayed.
persist <static dynamic>	filters by persistence.
type <ingress egress transit>	filters by tunnel type.
state <up down>	filters by operational state.
source <IP address>	filters by source IP address.
destination <IP address>	filters by destination IP address.
next-hop <IP address>	filters by next hop IP address.
out-ip-intf <Interface Supports Signals>	filters by outgoing IP interface.
in-label <NUMBER: 16- 1044479>	filters by inbound label.
out-label <NUMBER: 16- 1044479>	filters by outbound label.
recovery <protected unprotected>	filters by LSP recovery type.
role <primary backup locally- repaired active- backup>	filters by LSP protection role.
cos-profile <MPLS Tunnel COS Profile>	filters by tunnel CoS profile name

To display TP tunnels filtered by attribute

4 Display TP tunnels filtered by attribute:

```
gmpls tp-tunnel show matching-lsp <matching-lsp> {persist
<static|dynamic>} {type <ingress|egress|transit>} [path-
type <corouted | unidirectional>] [state <up|down>]
[source <IP address>] [destination <IP address>] [next-
```



```

hop <IP address>] [prev-hop <IP address>] [fwd-out-ip-
intf <signal-interface-object>] [rev-out-ip-intf
<Interface Supports Data>] [fwd-out-ip-intf <Interface
Supports Data>] [fwd-in-label <NUMBER: 16-1044479>] [fwd-
out-label <NUMBER: 16-1044479>] [rev-in-label <NUMBER:
16-1044479>] [rev-out-label <NUMBER: 16-1044479>]
[recovery <protected|unprotected>] [role
<primary|backup|locally-repaired|active-backup>] [cos-
profile <MPLS Tunnel COS Profile>] [class-type <NUMBER:
0-7>] [auto-size <enable>] [auto-size-interval <MINUTES>]
[auto-size-failure <alarm | mbb>] [min-bandwidth <Kbps>]
[max-bandwidth <Kbps>] [increment-bandwidth <Kbps>] [lsp-
reopt <enable>] [oper-bandwidth <Kbps>]

```

where

matching-lsp <matching-lsp>	is the tunnel to be displayed.
persist <static dynamic	filters by persistence.
type <ingress egress transit>]	filters by tunnel type.
path-type <corouted unidirectional>]	filters by tunnel path type.
state <up down>	filters by operational state.
source <IP address>	filters by source IP address.
destination <IP address>	filters by destination IP address.
next-hop <IP address>	filters by next hop IP address.
prev-hop- <IP address>	filters by previous hop IP address.
rev-out-ip-intf <Interface Supports Data>	filters by reverse outgoing IP interface.
fwd-out-ip-intf <Interface Supports Data>]	filters by forward outgoing IP interface.
fwd-in-label <NUMBER: 16- 1044479>	filters by forward in-bound label.
fwd-out-label <NUMBER: 16- 1044479>	filters by forward out-bound label.

where

rev-in-label <NUMBER: 16-1044479>]	filters by reverse in-bound label.
rev-out-label <NUMBER: 16-1044479>	filters by reverse out-bound label.
recovery <protected unprotected>	filters by failure recovery type.
role <primary backup locally-repaired active-backup>	filters by protection role.
cos-profile <MPLS Tunnel COS Profile>	filters by tunnel CoS profile name
class-type <NUMBER: 0-7>	filters by class type
auto-size <enable>	filters by auto sizing
auto-size-interval <MINUTES>	filters by auto size interval
auto-size-failure <alarm mbb>	filters by auto sizing failure handling mode
min-bandwidth <Kbps>	filters by minimum bandwidth (Kbps)
max-bandwidth <Kbps>	filters by maximum bandwidth (Kbps)
increment-bandwidth <Kbps>	filters by increment bandwidth (Kbps)
lsp-reopt <enable>	filters by LSP re-optimization
oper-bandwidth <Kbps>	if zero, then filters for exact match, else filters for operational bandwidth greater than or equal to the specified bandwidth

To display dynamic ingress TE tunnels

- 5 Display dynamic ingress TE tunnels:

```
mpls tunnel show rsvp-ingress <rsvp-ingress>
```

To display dynamic transit TE tunnels

- 6 Display dynamic transit TE tunnels:

```
mpls tunnel show rsvp-transit <rsvp-transit>
```

To display traffic statistics for dynamic transit TE-tunnels

- 7 Display traffic statistics for dynamic transit TE-tunnels:
 mpls tunnel show rsvp-transit <rsvp-transit> statistics

To display dynamic co-routed tunnels

- 8 Display dynamic co-routed tunnels:
 gmpls tp-tunnel show rsvp-ingress-corout <rsvp-ingress-corout>

To display dynamic unidirectional tunnels

- 9 Display dynamic unidirectional tunnels:
 gmpls tp-tunnel show rsvp-ingress-unidir <rsvp-ingress-unidir>

—end—

Example

This is a sample display of traffic statistics for a dynamic transit TE-tunnel:

```
5142-133*> mpls tunnel show rsvp-transit TE_dyn_uni_199_140 statistics
```

-----MPLS TRANSIT TUNNEL TRAFFIC STATISTICS-----	
Parameter	Value
Tunnel Index	786433
Tunnel Name	TE_dyn_uni_199_140
Rx Packets	13268421
Tx Packets	13268590
Rx Bytes	1246838774
Tx Bytes	1246854660

Procedure 9-12

Displaying auto-sizing statistics

Display auto-sizing statistics to learn more about how auto-sizing is used to adjust bandwidth. [Table 9-7](#) describes the fields in the PM Threshold Interval Profile table.

Table 9-7
Fields in the PM Threshold Interval Profile table

Field	Description
Sample Interval	Sample interval is configured by the user for an LSP. At the sample interval, bandwidth utilization statistics are collected for the LSP and recorded.
Maximum Average Bandwidth	Amongst a collection of bandwidth utilization statistics, the maximum average bandwidth is the one that has the highest value.
Adjustment Threshold	User-specified value is a percentage or actual value in MB Formula is DIFF (Maximum Average Bandwidth, Current LSP BW) exceeds Adjustment-Threshold. The Adjustment Threshold is applied to underutilization and overutilization detection.
Adjustment Interval	User configurable time interval in minutes System performs the Adjustment Threshold crossing check at the end of the timer interval Several multiples of sample interval value. If the sample interval is, for example, 120 seconds, the Adjustment Interval value must be a multiple of 2 minutes
Overflow Utilization Threshold	This user configurable value is a percentage or an absolute value that can be the same as the Adjustment Threshold or different. The formula is DIFF (Bandwidth Utilization, current LSP BW) > Overflow Utilization Threshold. The Overflow Utilization Threshold is checked at every Sample Interval.
Overflow Utilization Count	When the Overflow Utilization Threshold is exceeded, Overflow Utilization Count is incremented. If the consecutive count exceeds the user configured Overflow Utilization Count, then the Overflow is detected and the maximum bandwidth value during the count is used for LSP upsizing. The overflow detection is as rapid as the count * sample-interval.
Underflow Utilization Threshold	User-configured value in percentage or absolute value. This value can be the same as the Adjustment Threshold or different. The formula to determine if the threshold is exceeded is: DIFF (Bandwidth Utilization, current LSP BW) > Underflow Utilization Threshold. Underflow Utilization Threshold is checked at every Sample Interval. For example, if Underflow Utilization Threshold is 10% and LSP B/W is 10MB then Bandwidth Utilization less than 9MB then Underflow Utilization Threshold is exceeded.

Table 9-7
Fields in the PM Threshold Interval Profile table

Field	Description
Underflow Utilization Count	Underflow Utilization Count is incremented whenever Underflow Utilization Threshold is crossed. If the consecutive count exceeds the user-configured Underflow Utilization Count, then underflow is detected and the maximum bandwidth value during the count is used for LSP downsizing. The underflow detection is as rapid as count * sample-interval. Users can slow down the downsizing by configuring the underflow count to a high value. This helps the scale and frequent upsizing and downsizing

Step	Action

1	Display auto-sizing statistics: <pre>pm threshold show interval-profile <interval-profile></pre> —end—

Example

The following example shows sample output for the pm threshold show interval-profile command for an interval-profile named AutoSize.

```
pm threshold show interval-profile AutoSize
```

PM THRESHOLD INTERVAL PROFILE	
Parameter	Value
Interval Profile Name	AutoSize
Interval Profile ID	1
Sample Interval	5 min
Adjustment Interval	10m
Adjustment Threshold	2
Adjustment Threshold Percentage	2
Overflow Threshold	2
Overflow Threshold Percentage	2
Overflow Count	2
Underflow Threshold	2
Underflow Threshold Percentage	2
Underflow Count	2
Alert Interval	5m
Alert Threshold Percentage	

PM THRESHOLD INTERVAL PROFILE USAGE	
PM Instance	Instance Oper State
AutoSizeInstance	enabled

Procedure 9-13

Configuring RSVP-TE graceful restart helper mode

Enable global supports graceful helper mode for RSVP-TE.

Graceful restart can only be activated on newly-created RSVP tunnels.

Step	Action
1	Enable graceful restart helper mode for RSVP-TE: <code>rsvp-te graceful-restart enable</code>
2	Configure the graceful restart helper RSVP-TE timer: <code>rsvp-te set ip-interface <ip-interface> hello-interval <SECONDS: 0..30></code>
3	(Optional) Display RSVP-TE timers: <code>rsvp-te graceful-restart enable</code>
—end—	

Example

The following example shows sample output for the `rsvp-te show timers` command:

```
> rsvp-te show timers
+-----RSVP TE PROTOCOL TIMERS-----+
| Name                                | Value                                |
+-----+-----+
| Retry Interval (sec)                 | 3                                    |
| Restart Time (sec)                   | 60                                   |
| Tunnel Recovery Time (sec)            | 0                                    |
| PW Recovery Time (sec)                | 0                                    |
| Retry Max Default                     | 10                                   |
| Retry Infinite State                  | On                                   |
| Activation Interval (sec)              | 300                                  |
| Total Activation Period (sec)          | Infinite                             |
| Pseudowire Capability                 | On                                   |
| Pseudowire Status                     | On                                   |
+-----+-----+
```

Procedure 9-14

Configuring T-LDP graceful restart helper mode

Configure T-LDP graceful restart helper mode.

LDP Graceful Restart will be activated only on newly created LDP session

Step	Action

- 1 Enable graceful restart helper mode for T-LDP:

```
ldp graceful-restart enable
```

- 2 Set the graceful restart helper LDP timer:

```
ldp graceful-restart set max-peer-reconnect <NUMBER: 0..600>
```

where

max-peer-reconnect <NUMBER: 0..600>	Specifies the time that the sender of the TLV would like the receiver of that TLV to wait after the receiver detects the failure of an LDP session.

- 3 Set the graceful restart helper LDP recovery timer:

```
ldp graceful-restart set max-peer-recover <NUMBER: 0..600>
```

where

max-peer-recover <NUMBER: 0..600>	Specifies how long the device should hold stale label-FEC bindings after an RSVP session has been reestablished.

- 4 (Optional) Display LDP timers:

```
ldp graceful-restart enable
```

—end—

Example

The following example shows sample output for the `ldp show timers` command:

```
> ldp show timers
```

LDP Timers	
Parameter	Value
LDP Hello Hold Time (sec)	180
LDP Keepalive Hold Time (sec)	180
LDP GR Reconnect Time (sec)	0
LDP GR Recovery Time (sec)	0

LDP GR Max Peer Reconnect (sec)	180
LDP GR Max Peer Recovery (sec)	240
+-----+-----+	

Procedure 9-15

Configuring MPLS fast reroute

Configure the global attributes for MPLS fast reroute.

Step	Action

1	Configure the attributes for MPLS fast reroute:

```
mpls frr set {[global-reversion <enable|disable>]
[global-reversion-hold-time <NUMBER: 0..3600>] [global-
reversion-retry-interval <NUMBER: 30..3600>] [global-
reversion-retry-count <NUMBER: 0..100>] [facility-
bypass-autogeneration <enable|disable>] [plr-support
<enable|disable>] [auto-fb-cleanup-hold-time <NUMBER:
0..3600>]}
```

where

global-reversion <enable disable>	enables or disables global reversion of fast reroute. The default is enable.

global-reversion-hold-time <NUMBER: 0..3600>	sets the global reversion hold time, in seconds, of fast reroute. The default is 0.

[global-reversion-retry-interval <NUMBER: 30..3600>]	sets the global reversion retry interval, in seconds. The default is 60.

[global-reversion-retry-count <NUMBER: 0..100>]	sets the number of global reversion retries. The default is 1.

facility-bypass-autogeneration <enable disable>	enables or disables the facility bypass auto-generation of fast reroute. The default is disable.

plr-support <enable disable>	enables or disables PLR support of fast reroute. The default is enable.

auto-fb-cleanup-hold-time <NUMBER: 0..3600>	is the auto-generated facility bypass deletion hold time, in seconds, of fast reroute. The default is 120.

—end—

Procedure 9-16

Returning MPLS fast reroute to default values

Return the attributes for MPLS fast reroute to default values.

Step	Action

1	Return the attributes for MPLS fast reroute to default values: <pre>mpls frr unset {[global-reversion] [global-reversion- hold-time] global-reversion-retry-interval] [global- reversion-retry-count] [facility-bypass-autogeneration] [plr-support] [auto-fb-cleanup-hold-time]}</pre> —end—

Procedure 9-17

Creating an auto facility bypass profile

The properties of the auto-generated bypass tunnel are determined based on the bypass profile configured on the PLR. There can only be one bypass profile per interface. The bypass profile is used to determine the properties with which the bypass tunnel is created. Any updates in the bypass profile are not applied to the existing bypass tunnel and are used on new bypass tunnels created after profile updates.

The bypass tunnel is auto-deleted if no one uses it. The system waits for the auto facility bypass hold time before deletion.

While the auto facility bypass can be explicitly deleted, it cannot be disabled. Whether the auto facility bypass is manually or automatically deleted, an SNMP trap is generated.

While the auto facility bypass profile can be modified by the user at run-time, only new auto facility bypasses use the updated profile.

Step	Action

1	Create an MPLS auto facility bypass profile:

```
mpls tunnel-auto-fb-profile create auto-fb-profile <auto-
fb-profile-name string (1..31> {[bandwidth-percentage-
source <user | fast-reroute-object>] [bandwidth-
percentage-value <NUMBER: 0...100>] [setup-hold-
priority-source <user | fast-reroute-object>] [setup-hold-
priority-value <NUMBER: 0...7>] [link-preference <yes |
no>] [affinity-disjoint <yes | no>] [affinity-source
<user | fast-reroute-object | protected-interface-
disjoint> <resource-include-all <MPLS TE Admin Color
Group> | resource-include-any <MPLS TE Admin Color Group>
| resource-exclude-any <MPLS TE Admin Color Group> >]
[srlg-mode < maximal | strict | none>] [share-srlg-link
```

```
<RSVP SRLG values and ranges: 0..4294967295[-
<range>][,...]>] [share-srlg-node <RSVP SRLG values and
ranges: 0..4294967295[-<range>][,...]>]]}
```

where

auto-fb-profile <auto-fb-profile-name string (1..31>	is the auto facility bypass profile name.
[bandwidth-percentage-source <user fast-reroute-object>]	is the source from where the bandwidth percentage is selected. This value is user provided or is taken from FAST_REROUTE of the protected LSP. The default is user.
[bandwidth-percentage-value <NUMBER: 0...100>]	is the protected port bandwidth's percentage. This parameter is used if the bandwidth source is the user. The default is 0.
[setup-hold-priority-source <user fast-reroute-object>]	is the source from where the setup and hold priority is selected. This value is user provided or is taken from FAST_REROUTE of the protected LSP. This parameter is used if setup and hold priority source is the user. The default is user.
setup-hold-priority-value <NUMBER: 0...7>]	This parameter is used if setup and hold priority source is the user. The default is 0.
[link-preference <yes no>]	is the link preference for the auto facility bypass tunnel. The default is no.
[affinity-disjoint <yes no>]	is the resource affinity preference for the auto facility bypass tunnel. The default is yes.
[affinity-source <user fast-reroute-object protected-interface-disjoint> <resource-include-any <MPLS TE Admin Color Group> <resource-include-all <MPLS TE Admin Color Group> <resource-exclude-any <MPLS TE Admin Color Group>>]	is the source from where resource affinity is selected. This parameter is user provided or is taken from FAST_REROUTE of the protected LSP or is taken from the protected interface's affinity. The default is fast-reroute-object. Used only if "affinity-disjoint" is set. There is no default MPLS TE Admin Color Group in the system. Includes any, includes all, or excludes any MPLS TE Admin Color Groups.

where

[srlg-mode < maximal | strict | none>]

SRLG diversity is related to the protected LSP links between the PLR and merge point. The bypass tunnel should be SRLG disjoint if this flag is enabled. The default is none.

[share-srlg-link] <RSVP SRLG values and ranges:

0..4294967295 [-<range>][,...]>]

is the shared SRLG to be “relaxed” from disjoint computation when the auto facility bypass provides link protection. The default is none (0,0,0,0) with four possible values.

[share-srlg-node] <RSVP SRLG values and ranges:

0..4294967295 [-<range>][,...]>]

is the shared SRLG to be “relaxed” from disjoint computation when the auto facility bypass provides node protection. The default is none (0,0,0,0) with four possible values.

—end—

Procedure 9-18

Configuring an auto facility bypass profile

Configure an auto facility bypass profile.

Step	Action

1	Configure an auto facility bypass profile: <pre>mpls tunnel-auto-fb-profile set fb-profile <auto-fb- profile-name string (1...31)> {[bandwidth-percentage- source <user fast-reroute-object> bandwidth-percentage- value <NUMBER: 1...100>] [setup-hold-priority-source <user fast-reroute-object> setup-hold-priority-value <NUMBER: 0...7>][link-preference <yes no>] [affinity- disjoint <yes no>] [affinity-source <user fast- reroute-object protected-interface-disjoint> <resource-include-all <MPLS TE Admin Color Group> resource-include-any <MPLS TE Admin Color Group> resource-exclude-any <MPLS TE Admin Color Group> >] [srlg-mode < maximal strict none>] [share-srlg-link <RSVP SRLG values and ranges: 0..4294967295[- <range>][,...]>] [share-srlg-node <RSVP SRLG values and ranges: 0..4294967295[-<range>][,...]>] [force-apply]}</pre> —end—

Procedure 9-19

Returning an auto facility bypass profile to default values

Return an auto facility bypass profile to default values.

Step	Action

1	Return an auto facility bypass profile to default values: <pre>mpls tunnel-auto-fb-profile unset fb-profile <auto-fb-profile-name string (1...31)> {[bandwidth-percentage-source] [bandwidth-percentage-value] [setup-hold-priority-source] [setup-hold-priority-value] [affinity-disjoint <yes no>] [affinity-source <resource-include-all resource-include-any resource-exclude-any> [srlg-mode] [share-srlg-link] [share-srlg-node] [link-preference]</pre>

—end—

Procedure 9-20

Deleting an auto facility bypass profile

Delete an auto facility bypass profile.

Step	Action

1	Delete an auto facility bypass profile. <pre>mpls tunnel-auto-fb-profile delete auto-fb-profile <auto-fb-profile-name string (1..31)></pre> where auto-fb-profile <auto-fb-profile-name string (1..31)> is the name of the auto facility bypass profile that you are deleting. —end—

Procedure 9-21

Creating a fast reroute profile

Create a fast reroute profile that can be shared by the multiple fast reroute tunnels. The fast reroute profile defines the parameters signaled in FAST_REROUTE object while signaling the fast reroute tunnels.

Configuring the fast reroute profile associates the profile to the LSP protection is desired for.

Step	Action

1	Create a fast reroute profile:

```
mpls tunnel-frr-profile create frr-profile <frr-profile-
name string (1...31)> {[bw-protection <yes | no>] [node-
protection <yes | no>] [bandwidth <Kbps:
1000..1000000000>] [setup-priority <NUMBER: 0..7>] [hold-
priority <NUMBER: 0..7>] [resource-include-any <MPLS TE
Admin Color Group>] [resource-include-all <MPLS TE Admin
Color Group>] [resource-exclude-any <MPLS TE Admin Color
Group>]}
```

where

<frr-profile-name string (1...31)> is the tunnel profile name.

[bw-protection <yes | no>] enables or disables bandwidth protection. The default is no.

[node-protection <yes | no>] enables or disables node protection. The default is yes.

[bandwidth <Kbps: 1000..1000000000>] is the bandwidth in Kbps. The default is 1000.

[setup-priority <NUMBER: 0..7>] is the fast reroute tunnel setup priority. The default is 0.

[hold-priority <NUMBER: 0..7>] is the fast reroute tunnel hold priority. The default is 0.

[resource-include-any <MPLS TE Admin Color Group>] includes any MPLS TE Admin Color group. There is no default color group in the system.

[resource-include-all <MPLS TE Admin Color Group>] includes all MPLS TE Admin Color groups. There is no default color group in the system.

[resource-exclude-any <MPLS TE Admin Color Group>] excludes any MPLS TE Admin Color Group. There is no default color group in the system.

—end—

Procedure 9-22

Configuring a fast reroute profile

Configure a fast reroute profile.

Step	Action

1	Configure a fast reroute profile:

```
mpls tunnel-frr-profile set frr-profile <frr-profile-name
string (1...31)> {[bw-protection <yes | no>] [node-
protection <yes | no>] [bandwidth <Kbps:
1000..1000000000>] [setup-priority <NUMBER: 0..7>] [hold-
priority <NUMBER: 0..7>] [resource-include-any <MPLS TE
Admin Color Group>] [resource-include-all <MPLS TE Admin
Color Group>] [resource-exclude-any <MPLS TE Admin Color
Group>]}
```

where

<frr-profile-name string (1...31)> is the tunnel profile name.

[bw-protection <yes | no>] enables or disables bandwidth protection. The default is no.

[node-protection <yes | no>] enables or disables node protection. The default is yes.

[bandwidth <Kbps: 1000..1000000000>] is the bandwidth in Kbps. The default is 1000.

[setup-priority <NUMBER: 0..7>] is the fast reroute tunnel setup priority. The default is 0.

[hold-priority <NUMBER: 0..7>] is the fast reroute tunnel hold priority. The default is 0.

[resource-include-any <MPLS TE Admin Color Group>] includes any MPLS TE Admin Color group. There is no default color group in the system.

[resource-include-all <MPLS TE Admin Color Group>] includes all MPLS TE Admin Color groups. There is no default color group in the system.

[resource-exclude-any <MPLS TE Admin Color Group>] excludes any MPLS TE Admin Color Group. There is no default color group in the system.

—end—

Procedure 9-23

Returning the fast reroute profile to default values

Return the fast reroute profile to default values.

Step	Action

1	Return a fast reroute profile to default values: <pre>mpls tunnel-frr-profile unset frr-profile <frr-profile- name string (1..31)> {[bw-protection] [node-protection] [bandwidth] [setup-priority] [hold-priority] [color- group-include-any] [color-group-include-all] [color- group-exclude-any]}</pre> —end—

Procedure 9-24

Deleting a fast reroute profile

Delete a fast reroute profile.

Step	Action

1	Delete a fast reroute profile: <pre>mpls tunnel-frr-profile delete frr-profile <frr-profile- name string (1..31)> —end—</pre>

Procedure 9-25

Creating a fast reroute tunnel

Create a fast reroute tunnel.

Step	Action

1	Create a fast reroute profile:

<pre>mpls tunnel create rsvp-ingress <tunnel-name string (1..31)> {[frr-type <protected facility-bypass>] [frr- profile <frr-profile-name string (1..31)>]}</pre>

where

rsvp-ingress
<tunnel-name
string (1..31)>

is the rsvp-te tunnel name.

frr-type
<protected |
facility-bypass>

frr-type is optional for a UNI tunnel. It is required for FRR protected or facility bypass tunnels. An frr-type protected tunnel cannot be unset/set on an admin-enabled tunnel.

frr-profile <frr-
profile-name
string (1..31)>

is the frr-profile name.

—end—

Example

The following example creates a fast reroute tunnel with a default FRR profile.

```
> mpls tunnel create rsvp-ingress frr_protected_tnl1 dest-ip 5.5.5.5 frr-type
protected
```

The following example creates a fast reroute tunnel with a user-defined FRR profile.

```
> mpls tunnel create rsvp-ingress dest-ip 5.5.5.5 frr-type protected_tnl1 frr-
type protected frr-profile fprofil
```

Procedure 9-26

Creating a facility bypass tunnel

The facility bypass tunnels can be created on a PLR or can be auto-generated by the PLR during the creation of the backup paths. This procedure describes how the user can create a facility bypass tunnel.

Step	Action
1	Create a facility bypass tunnel: <pre>mpls tunnel create rsvp-ingress <tunnel-name string (1..31)> {[frr-type <protected facility-bypass>] [protected-interface <interface name string (1..31)>] [explicit-tunnel-path <MPLS Rsvp Path>] [frr-exclude-ip <IP address>] [srlg-mode <maximal strict none>] [share-srlg <RSVP SRLG values and ranges: 0..4294967295[-<range>] [, ...]>]}</pre> where rsvp-ingress <tunnel-name string (1..31)> is the rsvp-te tunnel name. For facility bypass tunnel, the tunnel name must end with <code>_facility</code>. frr-type <protected facility-bypass> frr-type is optional for a UNI tunnel. It is required for FRR protected or facility bypass tunnels. The tunnel name must contain <code>_facility</code>. For an frr-type other than facility-bypass, <code>_facility</code> in the tunnel name is invalid. The protected-interface is MUST and it cannot be modified later. Advanced MPLS features do not apply to facility bypass tunnel. The setup/hold priority is always zero for a facility bypass tunnel. If a user configures this value to a non-zero value, it is internally changed to zero. protected-interface <interface <interface name string (1..31)> is the protected-interface name. explicit-tunnel-path <MPLS Rsvp Path> This value can be can either provide explicit-tunnel path or configure frr-exclude-ip (and srlg). If this value is provided, facility-bypass tunnel is created using the provided ERO and there shall not be any additional SRLG considerations. Explicit-path cannot be configured by the user with exclude-ip and srlg and vice versa.

where

frr-exclude-ip <IP address> If this value is included, that hop is excluded from the tunnel path. The tunnel is dynamically routed using a path that does not contain exclude-ip. In this case, you can specify the srlg-disjoint mode to make the tunnel srlg disjoint of the protected interface.

srlg-mode <maximal | strict | none> If this value is configured the facility bypass tunnel is created as the SRLG disjoint of the protected interface.

share-srlg <RSVP SRLG values and ranges: 0..4294967295[-<range>][,...]> is the shared SRLG. The default is none (0,0,0,0) with four possible values.

—end—

Example

The following example creates a fast reroute tunnel with an explicit-tunnel-path.

```
> mpls tunnel create rvsp-ingress frr_tnl2_facility dest-ip 4.4.4.4 frr-type
facility-bypass protected-interface intf1 explicit-tunnel-path backup
```

The following example creates a fast reroute facility-bypass tunnel with bandwidth and frr-exclude-ip.

```
> mpls tunnel create rsvp-ingress frr_tunl2_facility dest-ip 4.4.4.4 frr-type
facility-bypass protected-interface intf1 min-bandwidth 6000 max-bandwidth
6000 frr-exclude-ip 110.1.1.2 srlg-mode strict share-srlg 2001,2002,2003,2004
```


Procedure 9-27

Configuring a fast reroute tunnel

MPLS-TE tunnels are configured on the ingress LER with frr protection enabled and a default or user-created fast reroute profile.

Step	Action

1	Create a fast reroute tunnel:

```
mpls tunnel set rsvp-ingress <tunnel-name string
(1..31)> {[frr-type <protected | facility-bypass>] [frr-
profile <frr-profile-name string (1..31)>] [protected-
interface <interface name string (1..31)>]}
```

where

rsvp-ingress is the rsvp-te tunnel name.
 <tunnel-name
 string (1..31)>

frr-type frr-type is optional for a UNI tunnel. It is required for FRR
 <protected | protected or facility bypass tunnels.
 facility-bypass> The tunnel name must contain _facility. For an frr-type other
 than facility-bypass, _facility in the tunnel name is invalid.
 The protected-interface is MUST and it cannot be modified
 later.

Advanced MPLS features do not apply to facility bypass
 tunnel.

The setup/hold priority is always zero for a facility bypass
 tunnel. If a user configures this value to a non-zero value, it
 is internally changed to zero.

protected-
 interface is the protected-interface name.
 <interface
 <interface name
 string (1..31)>

—end—

Procedure 9-28

Returning fast reroute tunnel attributes to default values

Return fast reroute tunnel attributes to default values.

Step	Action

1	Return fast reroute tunnel attributes to default values: <pre>mpls tunnel unset rsvp-ingress <tunnel-name string (1...31)> {[frr-type] [frr-profile] [protected- interface]}</pre>

—end—

Procedure 9-29

Deleting a fast reroute tunnel

Delete a fast reroute tunnel.

Step	Action

1	Delete a fast reroute tunnel: <pre>mpls tunnel delete rsvp-ingress <tunnel-name string (1..31)></pre>

—end—

Procedure 9-30

Configuring the IP interface with TE attributes

Configure the IP interface with TE attributes.

Step	Action

1	Configure the IP interface with different TE attributes: <pre>mpls traffic-eng set ip-interface <ip-interface-name string (1..31)> {[frr-auto-fb-profile <auto-fb-profile- name string (1..31)>]</pre> where ip-interface <ip- interface-name string (1..31)> is the IP interface name. frr-auto-fb-profile is the MPLS auto-generated facility bypass profile <auto-fb-profile-name name. string (1..31)>

—end—

Procedure 9-31

Returning the IP interface TE attributes to default values

Return the IP interface with TE attributes to default values.

Step	Action

1	Return the IP interface TE attributes to default values: <pre>mpls traffic-eng unset ip-interface <ip-interface-name> string (1..31)> {[frr-auto-fb-profile]} —end—</pre>

Procedure 9-32

Displaying MPLS fast reroute values

Display MPLS fast reroute values.

Step	Action

1	Display MPLS fast reroute values:

	<code>mpls frr show</code>

	<code>—end—</code>

Example

The following example shows the output from the `mpls frr show` command.

```
> mpls frr show
```

-----MPLS FRR GLOBAL-----	
Parameter	Value
PLR Support	Enabled
Global Reversion	Enabled
Global Reversion Hold Time (sec)	0
FB Auto Generation	Disabled
Auto-Fb Cleanup Hold Time	120
Global Reversion Retry Interval	60
Global Reversion Retry Count	1

Procedure 9-33

Displaying auto facility bypass profiles

You can display

- all auto facility bypass profiles
- a specific facility bypass profile

Step	Action

To display all auto facility bypass profiles

- 1 Display all auto facility bypass profiles:

```
mpls tunnel-auto-fb-profile show
```

To display a specific auto facility bypass profile

- 2 Display a specific auto facility bypass profile:

```
mpls tunnel-auto-fb-profile show auto-fb-profile <auto-  
fb-profilename string (1..31)>
```

—end—

Example

The following example shows the output from the mpls tunnel-auto-fb-profile show command.

```
> mpls tunnel-auto-fb-profile show
```

```
+---- MPLS TUNNEL AUTO FB PROFILE TABLE-----+
| Profile |      Profile Name      | UseCnt |
| Index   |                         |        |
+-----+-----+-----+
| 1       | Default                | 3      |
| 2       | SampleFB1              | 2      |
+-----+-----+-----+
```

The following example shows the output for the auto-fb-profile prof1.

```
> mpls tunnel-auto-fb-profile show auto-fb-profile prof1
```

```
+-----MPLS TUNNEL AUTO FB PROFILE DETAILS-----+
| Parameter | Value |
+-----+-----+
| Profile Name | prof1 |
| Profile index | 2     |
| Profile use count | 2     |
| Protected Interface | ip40, ip110 |
| Setup/Hold Priority Source | User |
| Setup/Hold Priority | 0     |
| Bandwidth Source | User |
| Bandwidth Percentage | 0     |
| Link Preference | NO    |
+-----+-----+
```

Affinity Disjoint	YES
Affinity Source	FAST_REROUTE Object
Resource Include-any	-
Resource Include-all	-
Resource Exclude-any	-
SRLG Mode	None
Share SRLG Node	10,20
Share SRLG Link	10,20,30

Procedure 9-34

Displaying fast reroute profiles

You can display

- all fast reroute profiles
- a specific fast reroute profile

Step	Action

To display all fast reroute profiles

- 1 Display all fast reroute profiles:

```
mpls tunnel-frr-profile show
```

To display a specific fast reroute profile

- 2 Display a specific fast reroute profile:

```
mpls tunnel-frr-profile show frr-profile <frr-profile-  
name string (1..31)>
```

—end—

Example

The following example shows the output from the mpls tunnel-frr-profile show command.

```
> mpls tunnel-frr-profile show
```

-----MPLS TUNNEL FAST-REROUTE PROFILE TABLE-----				
Profile Name	Protection BW	Node	Protection Method	UseCnt
DefaultFrrProfile	NO	YES	facility	0

The following example shows the output for the frr-profile fprof1.

```
> mpls tunnel-frr-profile show frr-profile fprof1
```

-----MPLS TUNNEL FAST-REROUTE PROFILE DETAILS-----	
Profile Name	fprof1
Profile index	2
Profile use count	0
Setup Priority	0
Hold Priority	0
Hop Limit	-
Bandwidth	1000
Bandwidth Protection	YES
Node Protection	YES
Protection Method	Facility
Color Group Include Any	RED
Color Group Include All	-

Color Group Exclude Any	-	
+-----+	+-----+	+-----+

Procedure 9-35

Displaying fast reroute protected tunnels at ingress

Display fast reroute protected tunnels at ingress.

Step	Action

1	Display fast reroute protected tunnels at ingress: <pre>mpls tunnel show rsvp-ingress <tunnel-name string (1..31)></pre>

—end—

Example

The following example shows the output from the mpls tunnel show rsvp-ingress command for the tunnel frr_protected_tn11.

```
> mpls tunnel show rsvp-ingress frr_protected_tn11
```

-----MPLS INGRESS TE-TUNNEL DETAILS-----	
Parameter	Value
Tunnel Name	AB1_node
Tunnel Index	2
Tunnel Type	Dynamic
Nodal Role	Ingress
Destination IP Address	3.3.3.3
Next-Hop IP Address	130.1.1.2
Tunnel Group Index	32770
Admin State	Enabled
Oper State	Enabled
Out-Label	32768
LSP ID	1
Path Name	primary_path
Setup Priority	0
Hold Priority	0
Record Route	On
CSPF Route Selection	On
Forward Protection Role	Primary
Forward Protection State	Active
Forward Backup Tunnel Name	None Present
LSP Reoptimization	Disabled
LSP Reoptimization Interval (minutes)	0
LSP Reoptimization Time Remaining	00m 00s
Tunnel Reversion	On
Reversion Hold-Time	30
Forward CoS Profile Name	DefaultTunlCoSProfile
Forward CoS Profile Index	1
TTL Policy	fixed
Fixed TTL	255
Queue Group Profile Name	Default-QueueGroupProfile
Queue Group Profile Index	0
Queue Group Instance Number	0
Frr Signalling	On
Frr Profile Name	fprofil
Frr Profile index	2
Frr Tunnel Type	FRR Protected

Frr Protecting FB's Vif	2001
Frr Local Protection InUse	No
Class Type	0
Minimum Bandwidth (Kbps)	0
Maximum Bandwidth (Kbps)	1000000000
Increment Bandwidth (Kbps)	0
Auto Sizing	Disabled
Auto Sizing Mode	CAC
Auto Sizing Interval (minutes)	0
Auto Sizing Time Remaining	00m 00s
Auto Sizing Failure Handler	alarm
Requested bandwidth (Kbps)	0
Used bandwidth (Kbps)	0
Operational bandwidth(Kbps)	0

-----MPLS Actual Route TABLE-----					
Index	Node Ip	Local Protection		Bandwidth	Node
		Available	In Use	Protected	Protected
1	130.1.1.1	YES	NO	NO	YES
2	130.1.1.2	YES	YES	NO	NO
3	20.1.1.2	NO	NO	NO	NO

-----MPLS FRR PROTECTED TUNNEL TABLE-----	
Parameter	Value
FRR Profile Name	fprof1
FRR Profile Index	2
Local Protection Avail	Yes
Local Protection FB	tunl1_facility
Local Protection InUse	Yes

-----FRR Initial Route TABLE-----	
Index	Node Ip
1	130.1.1.1
2	130.1.1.2
3	120.1.1.1
4	20.1.1.2

Procedure 9-36

Displaying fast reroute protected tunnels at transit

Display fast reroute protected tunnels at transit.

Step	Action

1	Display fast reroute protected tunnels at transit: <pre>mpls tunnel show rsvp-transit <tunnel-name string (1..31)></pre>

—end—

Example

The following example displays fast reroute tunnels at transit for the tunnel frr_protected_tn12.

```
mpls tunnel show rsvp-transit frr_protected_tn12
```

-----MPLS TRANSIT TE-TUNNEL DETAILS-----	
Parameter	Value
Tunnel Name	frr_protected_tn12
Tunnel Index	1
Transit Type	Dynamic
Nodal Role	Transit
Destination IP Address	2.2.2.2
Source IP Address	4.4.4.4
Next-Hop IP Address	50.1.1.1
Admin State	Enabled
Oper State	Enabled
Forward In-Label	8192
Forward Out-Label	8192
Forward CoS Profile Name	DefaultTunlCoSProfile
Forward CoS Profile Index	1
TTL Policy	decrement
Fixed TTL	0
Xc index	2
Prev-Hop IP Address	22.1.1.2
-----MPLS FRR PROTECTED TUNNEL TABLE-----	
Parameter	Value
Local Protection Avail	Yes
Local Protection FB	frr_fac_bypass_tn11
Local Protection InUse	No

Procedure 9-37

Displaying IP interface information with TE attributes

Display IP interface information with TE attributes.

Step	Action

1	Display IP interface information with TE attributes. <pre>mpls traffic-eng show ip-interface <ip-interface-name> string (1..31)></pre>

—end—

Example

The following example displays IP information with TE attributes.

```
> mpls tunnel show ip-interface ip66
```

-----Traffic Eng Interface-----								
Interface Name		ip66						
Interface Index		4						
TE Metric		1						
Color Group Index		0						
Color Group Name		None						
SRLGs		None						
Maximum BW(kbps)		10000000						
Maximum Reservable BW(kbps)		10000000						
Class Profile Index		0						
Class Profile Name		None						
Auto FB Profile Index		1						
Auto FB Profile Name		DefaultAutoFbProfile						
TE Mode		DiffServ						
Class Type	Queue Group	Total BW (kbps)	Reserved BW (kbps)	Unreserved BW (kbps)	LOM	LOM Used	Alarm (%)	
CT0		10000000	0	10000000	1	N	100	
CT1		0	0	0	1	N	100	
CT2		0	0	0	1	N	100	
CT3		0	0	0	1	N	100	
CT4		0	0	0	1	N	100	
CT5		0	0	0	1	N	100	
CT6		0	0	0	1	N	100	
CT7		0	0	0	1	N	100	

Procedure 9-38

Configuring the soft preemption timer on a node

Configure the soft preemption timer on a node.

Step	Action
1	Set the soft preemption timer on a node: <pre>mpls soft-preemption set timer <SECONDS: 5..60> [default 30]</pre>
2	Display the soft preemption timer set on a node: <pre>mpls soft-preemption show</pre> —end—

Procedure 9-39

Creating an MPLS-TE unidirectional dynamic LSP with soft preemption

Create an MPLS-TE unidirectional dynamic LSP with soft preemption.

Step	Action

1	Create an MPLS-TE unidirectional dynamic LSP with soft preemption: mpls tunnel create rsvp-ingress <tunnel-name> dest-ip <ip-addr> setup-priority <0..7> hold-priority <0..7> soft-preemption <on/off> —end—

Example

Here is an example of creating an MPLS-TE unidirectional dynamic LSP with soft preemption:

```
> mpls tunnel create rsvp-ingress lsp1 dest-ip 1.1.1.1 setup-priority 2 hold-  
priority 2 min-bandwidth 50000000 soft-preemption on
```

See [Procedure 9-40, “Displaying advertised bandwidth for a link”](#) on [page 9-114](#) for an example of the advertised bandwidth for the new LSP1.

Procedure 9-40

Displaying advertised bandwidth for a link

Display the advertised bandwidth for a link.

Step	Action

1	To display the advertised bandwidth for a link:

```
mpls debug tables show i3EtTbl
```

2	Alternatively, enter:

```
isis te-tables show te
```

—end—

Example

Here is sample bandwidth before creating a tunnel:

```
5160-PI-31*> mpls debug tables show i3EtTbl
```

-----I3 ET TABLE-----			
IfIndex		16	
IP Address		15.31.31.1	
TE Metric		5	
num_bandwidth_constraints		0	
bandwidth_constraint_model		0	

Unreserved BW	IEEE Floating Bytes	IEEE floating in hexa	kbps

max_bandwidth	1318388473	4e9502f9	10000000
max_resrv_bandwidth	1318388473	4e9502f9	10000000
unres_bwidth 0	1318388473	4e9502f9	10000000
unres_bwidth 1	1318388473	4e9502f9	10000000
unres_bwidth 2	1309999865	4e1502f9	10000000
unres_bwidth 3	1309999865	4e1502f9	10000000
unres_bwidth 4	1309999865	4e1502f9	10000000
unres_bwidth 5	1309999865	4e1502f9	10000000
unres_bwidth 6	1309999865	4e1502f9	10000000
unres_bwidth 7	1309999865	4e1502f9	10000000

Here is sample bandwidth after creating LSP1:

```
5160-PI-31*> mpls debug tables show i3EtTbl
```

-----I3 ET TABLE-----			
IfIndex		16	
IP Address		15.31.31.1	
TE Metric		5	
num_bandwidth_constraints		0	
bandwidth_constraint_model		0	

Unreserved BW	IEEE Floating Bytes	IEEE floating in hexa	kbps
max_bandwidth	1318388473	4e9502f9	10000000
max_resrv_bandwidth	1318388473	4e9502f9	10000000
unres_bwidth 0	1318388473	4e9502f9	10000000
unres_bwidth 1	1318388473	4e9502f9	10000000
unres_bwidth 2	1309999865	4e1502f9	5000000
unres_bwidth 3	1309999865	4e1502f9	5000000
unres_bwidth 4	1309999865	4e1502f9	5000000
unres_bwidth 5	1309999865	4e1502f9	5000000
unres_bwidth 6	1309999865	4e1502f9	5000000
unres_bwidth 7	1309999865	4e1502f9	5000000

Additional tunnel configuration

This chapter provides additional procedures for configuring tunnels on the 39XX/51XX Service Delivery, Aggregation and Virtualization Switches.

Ingress/egress CoS mapping

QoS is considered when the ingress PE maps the user priority bits in the VLAN header to the MPLS EXP/TC bits in the MPLS label stack. Transit label switch routers (LSRs) in the MPLS network can apply QoS policies to the Ethernet frames that are carried over pseudowires.


Incoming frames are mapped by a Resolved CoS Policy which defines how the internal R-CoS value is determined and is configured per port as follows:

- 1 Fixed (fixed-CoS) — CoS values in the frame are ignored and fixed R-CoS is applied to the frame from the Fixed Resolved CoS value for the port. When this policy is applied, CoS mapping is in untrusted mode. The default Fixed Resolved CoS value is 0.
- 2 Outer .1D mapped (dot 1d-tag1-cos) — The Priority Code Point (PCP)/ Layer 2 (L2) CoS 802.1D priority value from the outer-tag is mapped to an R-CoS value derived from the Resolved CoS Mapping table. When this policy is applied, CoS mapping is in trusted mode, which is the default for all ports.

CoS Mapping consists of the following steps:


- 1 The incoming frame is mapped by a Resolved CoS Policy using RCoSMap and assigned an R-CoS value defined by the CoS configuration for the respective port. The default Ingress FCoS -> RCoS Map: (RCoSMap) is Default FCoSRCoS. See [Figure 10-1](#).

Figure 10-1
Step 1: Ingress CoS -> RCoS Mapping (RCoS Map)



- 2 At egress, the frame assigned to R-CoS maps to an Egress Port Queue.
The default Ingress to Egress QMap: Default-RCoS. See [Figure 10-2](#).


Figure 10-2
Step 2: RCoS to Queue mapping



- 3 The outgoing frame cos values (1D cos, MPLS Tunnel TC, MPLS PW TC) are derived using FCoSMap configuration on the port. Egress RCoS -> FCoS Map: (FCoSMap) is DefaultRCoSFCoS. See [Figure 10-3](#).

Figure 10-3

Step 3: RCoS -> FCoS Mapping (FCoS Map)



[Figure 10-4](#) clarifies which RCoSMap and FCoSMap used at various ends in the packet modification in the MPLS network.

Figure 10-4

CoS mapping in an MPLS network




Table 10-1 clarifies which RCoSMap or FCoSMap is used at various ends in the packet forwarding in an MPLS network. Any RCoSMap or FCoSMap not specified in this table does not apply to the packet modification.

Table 10-1
Cos mapping applicability in an MPLS network

Node	Ingress and Egress Action	Fixed Policy	Mapped Policy
Ingress LER Node (AC --> PW)	AT AC apply Port RCoSMap	Supported	Supported
	AT PW apply Tunnel FcosMap and PW FCoSMap	Supported	Supported
Transit LSR Node	AT incoming tunnel, apply incoming Tunnel RCoSMap	Supported	Supported
	AT outgoing Tunnel, apply outgoing Tunnel FCoSMap	Supported	Supported
Egress LER Node (PW --> AC)	AT PW apply PW RCoSMap	Supported	Supported
	AT AC apply Port FCoSMap	Supported	Supported

Limitations

The following section describes the limitations that apply to ingress CoS:

- EPLS overrides everything on the port.
- EVPLS in port inherit dominates the other EVPLs on the port regardless of the port being in the same or a different virtual switch. For example, there can be mixed encap-cos-policy on the same parent port.

Port=5, vlan=400 encap-cos-policy = fixed

Port=5, vlan=500 encap-cos-policy = vs-inherit

Port=5, vlan=600 encap-cos-policy = port-inherit

All CVIDs (400, 500, 600) are applied to port-inherit policy. For CVID=400, 500, this is not the expected behavior. This is the limitation on all platforms. The T3 implementation is also broken when mixing some configurations on the same port. The behavior can depend on the order that the configuration is applied. This applies to Q-in-Q, PBT and MPLS

Another example is

Vs=VS1, Port=5, vlan=400 encap-cos-policy = fixed

Vs=VS2, Port=5, vlan=500 encap-cos-policy = vs-inherit

Vs=VS3, Port=5, vlan=600 encap-cos-policy = port-inherit

In this case, where there are multiple virtual switches sharing the same physical port, all EVPLs will have port-inherit as the encap-cos-policy. It is recommended that EVPLs with fixed and mapped encap-cos-policy are not mixed on the same physical port.

- Ingress PCP stamping is implemented using encap-cos-policy. For practical purposes, modifying the PCP bits in the customer payload and putting the payload in MPLS encapsulations is not useful. For end-2-end CoS, modifying TC bits in MPLS encapsulations and outer encap 802.1q/13 vlan PCP bits are sufficient. PCP bits in the customer payload (inside the encapsulation) are not changed for Tagged and Untagged-Data frames.

The following limitations apply to egress CoS:

- Egress CoS is set on a per physical port and not on a per virtual port. All CVID bundles on the parent physical port have the same egress CoS behavior. There is 1-to-1 mapping between the physical port and the MPLS AC (virtual port).
- Custom frame-cos-map on an egress port does not work.

See the following CoS procedures:

- [“Configuring untagged frames” on page 13-18](#)
- [“Configuring egress l2-transform CoS values” on page 13-20](#)
- [“Configuring CoS profiles for MPLS tunnels” on page 10-7](#)
- [“Configuring CoS profiles for MPLS pseudowire” on page 13-21](#)
- [“Displaying CoS profiles” on page 13-25](#)

MPLS EXP/TC bits

SAOS supports Differentiated Service Code Point (DSCP) or IP precedence and IEEE 802.1p CoS classifiers on the customer-edge interfaces of the ingress provider edge (PE) switch. DSCP or IP precedence classifiers are used for Layer 3 packets. IEEE 802.1p is used for Layer 2 packets.

When a packet enters a customer-edge interface of the ingress PE switch, the switch associates the packet with a CoS servicing level before putting it onto the label-switched path (LSP). The switches with the LSP use the CoS value set at the ingress PE switch. The CoS value in the packet is translated and encoded in the MPLS header EXP/TC bits.

The switch applies the default or the custom EXP/TC classifier and the default or the custom EXP rewrite rule to the MPLS-enabled interfaces. As rewrite rules affect only egress interfaces, the switch applies the EXP/TC rewrite rule only to those MPLS interfaces that are transmitting MPLS packets.

Procedures are:

- [“Configuring CoS profiles for MPLS tunnels” on page 10-7](#)
- [“Displaying a summary of CoS profiles for MPLS tunnels” on page 10-9](#)
- [“Moving co-routed tunnels to a new path” on page 10-10](#)

- [“Configuring a static co-routed primary tunnel with a dynamic associated tunnel as backup” on page 10-15](#)
- [“Configuring a dynamic associated tunnel as primary and a static co-routed tunnel as backup” on page 10-17](#)
- [“Configuring a primary tunnel with an explicit path and a backup tunnel without an explicit path” on page 10-19](#)
- [“Configuring a primary tunnel without an explicit path and a backup tunnel with an explicit path” on page 10-21](#)
- [“Configuring a primary tunnel and a backup tunnel without explicit paths” on page 10-23](#)

Procedure 10-1

Configuring CoS profiles for MPLS tunnels

Configure CoS profiles for MPLS tunnels.

For more information about CoS profiles, refer to “Quality of Service configuration,” in *SAOS 6.16 Advanced Ethernet Configuration*.

Step	Action														
1	Create an MPLS tunnel CoS profile: <pre>mpls tunnel-cos-profile create cos-profile <cos-profile> {logical-id <NUMBER>} [frame-cos-map <Frame Cos Map Name>] [frame-cos-policy <mapped fixed>] [fixed-tc <NUMBER: 0-7>] [rcos-profile <Resolved CoS Profile Name>] [resolved-cos-policy <mapped fixed>] [resolved-cos-fixed <NUMBER: 0-7>]</pre> where <table> <tr> <td>{logical-id <NUMBER>}</td><td>selects profile index.</td></tr> <tr> <td>frame-cos-map <Frame Cos Map Name>]</td><td>selects the frame CoS map.</td></tr> <tr> <td>frame-cos-policy <mapped fixed></td><td>selects the frame CoS policy.</td></tr> <tr> <td>fixed-tc <NUMBER: 0-7></td><td>selects the encapsulation fixed TC value.</td></tr> <tr> <td>rcos-profile <Resolved CoS Profile Name></td><td>selects the resolved-cos-profile.</td></tr> <tr> <td>resolved-cos-policy <mapped fixed>]</td><td>selects the resolved CoS policy.</td></tr> <tr> <td>[resolved-cos-fixed <NUMBER: 0-7>]</td><td>selects fixed resolved CoS value.</td></tr> </table>	{logical-id <NUMBER>}	selects profile index.	frame-cos-map <Frame Cos Map Name>]	selects the frame CoS map.	frame-cos-policy <mapped fixed>	selects the frame CoS policy.	fixed-tc <NUMBER: 0-7>	selects the encapsulation fixed TC value.	rcos-profile <Resolved CoS Profile Name>	selects the resolved-cos-profile.	resolved-cos-policy <mapped fixed>]	selects the resolved CoS policy.	[resolved-cos-fixed <NUMBER: 0-7>]	selects fixed resolved CoS value.
{logical-id <NUMBER>}	selects profile index.														
frame-cos-map <Frame Cos Map Name>]	selects the frame CoS map.														
frame-cos-policy <mapped fixed>	selects the frame CoS policy.														
fixed-tc <NUMBER: 0-7>	selects the encapsulation fixed TC value.														
rcos-profile <Resolved CoS Profile Name>	selects the resolved-cos-profile.														
resolved-cos-policy <mapped fixed>]	selects the resolved CoS policy.														
[resolved-cos-fixed <NUMBER: 0-7>]	selects fixed resolved CoS value.														
2	Confirm the configuration of the specific profile: <pre>mpls tunnel-cos-profile show cos-profile <cos-profile></pre>														
3	Display a summary of CoS profiles: <pre>mpls tunnel-cos-profile show cos-profile <cos-profile></pre> —end—														

Example

The following example creates an MPLS tunnel CoS profile named TE_TUN_COS_PROF.

```
mpls tunnel-cos-profile create cos-profile TE_TUN_COS_PROF frame-cos-policy
fixed fixed-tc 4 resolved-cos-policy fixed resolved-cos-fixed 6
```

The following sample output confirms the configuration.

```
mpls tunnel-cos-profile show cos-profile TE_TUN_COS_PROF
```

-----MPLS Tunnel-CoS-Profile Details-----	
Parameter	Value
CoS Profile Name	TE_TUN_COS_PROF
CoS Profile Index	2
Frame CoS Policy	fixed
Frame CoS Map Name	DefaultRcosFcos
Frame CoS Map ID	1
Fixed TC	4
Resolved CoS Policy	fixed
Resolved CoS Map Name	DefaultFcosRcos
Resolved CoS Map ID	1
Resolved CoS Fixed	6
Use Count	0

Procedure 10-2

Displaying a summary of CoS profiles for MPLS tunnels

Display a summary of CoS profiles for MPLS tunnels.

Step	Action

1	Display a summary of CoS profiles.

```
mpls tunnel-cos-profile show cos-profile <cos-profile>
```

—end—

Example

The following example shows sample output for the mpls tunnel-cos-profile show command.

```
mpls tunnel-cos-profile show
```

MPLS Tunnel-COS-Profile Table								
CoS-Mapping Profile Name	Index	FCoSPolicy	FCoSMapID	FixedTc	RCoSPolicy	RCoSMapID	RCoSFixed	UseCnt
DefaultTunlCoSProfile	1	mapped	1	0	mapped	1	0	8
TE_TUN_COS_PROF	2	fixed	1	4	fixed	1	6	0

Procedure 10-3

Moving co-routed tunnels to a new path

Move co-routed tunnel to a new path to reconfigure the primary tunnel over an alternate path while using the existing backup tunnel to continue carrying traffic.

The path for the backup tunnel can also be changed by disabling the backup tunnel administratively. When the path for the backup tunnel is changed, ensure that it is not carrying traffic.

Note: Transit tunnels are moved by deleting the existing transit tunnels and recreating transit tunnels to synchronize with the ingress/egress tunnel configuration.

Step	Action

Move primary tunnels

1	Disable primary tunnels: <pre>gmpls tp-tunnel disable static-ingress-corout <MplsIngressStaticTpTunl> gmpls tp-tunnel disable static-egress-corout <MplsEgressStaticTpTunl></pre> The system responds by moving traffic to the backup tunnels.
2	Modify the next hop and previous hop of the primary tunnels: <pre>gmpls tp-tunnel set static-ingress-corout <MplsIngressStaticTpTunl> next-hop-ip <IpAddress> gmpls tp-tunnel set static-egress-corout <MplsEgressStaticTpTunl> prev-hop-ip <IpAddress></pre>
3	Delete the existing transit tunnel: <pre>gmpls tp-tunnel delete static-transit-corout <MplsTransitStaticTpTunl></pre>
4	Create a transit tunnel with a new path: <pre>gmpls tp-tunnel create static-transit-corout <MplsTransitStaticTpTunl> dest-ip <IP Address> src-ip <IP Address> prev-hop-ip <IP Address> next-hop-ip <IP Address> reverse-out-label <NUMBER: 16-1044479> forward- in-label <NUMBER: 16-1044479> reverse-in-label <NUMBER: 16-1044479> forward-out-label <NUMBER: 16-1044479></pre>

5 Enable primary tunnels:

```
gmpls tp-tunnel enable static-ingress-corout  
<MplsIngressStaticTpTunl>
```

```
gmpls tp-tunnel enable static-egress-corout  
<MplsEgressStaticTpTunl>
```

The system responds by moving traffic to the primary tunnels after reversion.

Move backup tunnels**6 Disable backup tunnels:**

```
gmpls tp-tunnel disable static-ingress-corout  
<MplsIngressStaticTpTunl>
```

```
gmpls tp-tunnel disable static-egress-corout  
<MplsEgressStaticTpTunl>
```

7 Modify the next hop and previous hop of the backup tunnels:

```
gmpls tp-tunnel set static-ingress-corout next-hop-ip  
<IpAddress>
```

```
gmpls tp-tunnel set static-egress-corout prev-hop-ip  
<IpAddress>
```

8 Enable backup tunnels:

```
gmpls tp-tunnel enable static-ingress-corout  
<MplsIngressStaticTpTunl>
```

```
gmpls tp-tunnel enable static-egress-corout  
<MplsEgressStaticTpTunl>
```

—end—

Example

Figure 10-5 shows a sample tunnel configuration. The sections that follow describe how to create the ingress, transit, and egress tunnels and then how to move the primary tunnels while allowing traffic to flow on the backup tunnels.

Figure 10-5
Example tunnel configuration

Ingress 7.7.7.5	170.1.1.X	Transit 8.8.8.5	180.1.1.X	Egress 9.9.9.5
	Primary		Primary	
	70.1.1.X		80.1.1.X	
	Backup		Backup	
	7.8.14.X		11.8.14.X	
	New primary		New primary	

Create the tunnels

The following commands create the primary and backup ingress tunnels.

```
gmpls tp-tunnel create static-ingress-corout st_in_co_45_1
dest-ip 9.9.9.5 next-hop-ip 170.1.1.50 forward-out-label 31
reverse-in-label 41
```

```
gmpls tp-tunnel create static-ingress-corout st_in_co_bk_45
dest-ip 9.9.9.5 next-hop-ip 70.1.1.50 forward-out-label 50
reverse-in-label 60 backup-tunnel st_in_co_45_1
```

The following commands create the primary and backup transit tunnels.

```
gmpls tp-tunnel create static-transit-corout st_tr_co_pri_1
dest-ip 9.9.9.5 src-ip 7.7.7.5 prev-hop-ip 170.1.1.2 next-hop-
ip 180.1.1.2 reverse-out-label 41 forward-in-label 31 reverse-
in-label 42 forward-out-label 45
```

```
gmpls tp-tunnel create static-transit-corout st_tr_co_bkp_1
dest-ip 9.9.9.5 src-ip 7.7.7.5 prev-hop-ip 70.1.1.2 next-hop-ip
80.1.1.2 reverse-out-label 60 forward-in-label 50 reverse-in-
label 61 forward-out-label 51
```

The following commands create the primary and backup egress tunnels.

```
gmpls tp-tunnel create static-egress-corout st_eg_co_54_1 src-
ip 7.7.7.5 prev-hop-ip 180.1.1.50 reverse-out-label 42 forward-
in-label 45
```

```
gmpls tp-tunnel create static-egress-corout st_eg_co_54_bk_1
src-ip 7.7.7.5 prev-hop-ip 80.1.1.50 reverse-out-label 61
forward-in-label 51 backup-tunnel st_eg_co_54_1
```

Move the primary tunnels

The following commands disable the primary ingress and egress tunnels. The system moves the traffic to the backup tunnels.

```
gmpls tp-tunnel disable static-ingress-corout st_in_co_45_1
```

```
gmpls tp-tunnel disable static-egress-corout st_eg_co_54_1
```

The following commands modify the next hop and previous hop of the primary ingress and egress tunnels.

```
gmpls tp-tunnel set static-ingress-corout st_in_co_45_1 next-
hop-ip 7.8.14.50
```

```
gmpls tp-tunnel set static-egress-corout st_eg_co_54_1 prev-
hop-ip 11.8.14.50
```

The following commands delete the transit tunnel and recreate a transit tunnel with new next hop and previous hop values that synchronize with the ingress and egress tunnels.

```
gmpls tp-tunnel delete static-transit-corout st_tr_co_pri_1
```

```
gmpls tp-tunnel create static-transit-corout st_tr_co_pri_1
dest-ip 9.9.9.5 src-ip 7.7.7.5 prev-hop-ip 7.8.14.2 next-hop-ip
11.8.14.2 reverse-out-label 41 forward-in-label 31 reverse-in-
label 42 forward-out-label 45
```

Enable primary ingress and egress tunnels. The system moves the traffic to the primary tunnels after reversion.

```
gmpls tp-tunnel enable static-ingress-corout st_in_co_45_1
```

```
gmpls tp-tunnel enable static-egress-corout st_eg_co_54_1
```

Move the backup tunnels

The following commands disable the backup ingress and egress tunnels. If the backup tunnel was carrying traffic, the primary tunnel is active when the backup tunnel is disabled. The system switches the traffic over to the primary tunnel. The next hop and prev hop configuration of the backup tunnel can also be changed when it is in standby mode.

10-14 Additional tunnel configuration

```
gmpls tp-tunnel disable static-ingress-corout st_in_co_bk_45
gmpls tp-tunnel disable static-egress-corout st_eg_co_bk_54_1
```

The following commands modify the next hop and previous hop of the backup ingress and egress tunnels.

```
gmpls tp-tunnel set static-ingress-corout st_in_co_bk_45 next-
hop-ip <IpAddress>
gmpls tp-tunnel set static-egress-corout st_eg_co_bk_54_1 prev-
hop-ip <IpAddress>
```

Enable the backup ingress and egress tunnels.

```
gmpls tp-tunnel enable static-ingress-corout st_in_co_bk_45
gmpls tp-tunnel enable static-egress-corout st_eg_co_bk_54_1
```

Procedure 10-4

Configuring a static co-routed primary tunnel with a dynamic associated tunnel as backup

Configure a static co-routed primary tunnel with dynamic associated tunnel as backup if required for the network configuration.

Step	Action
1	Create a static ingress co-routed tunnel: <pre>gmpls tp-tunnel create static-ingress-corout <static-ingress-corout> dest-ip <IP Address> next-hop-ip <IP Address> forward-out-label <NUMBER: 16-1044479> reverse-in-label <NUMBER: 16-1044479>} bfd-monitor <enable disable></pre>
2	Create a dynamic ingress unidirectional tunnel with the static ingress co-routed tunnel configured as the backup tunnel: <pre>gmpls tp-tunnel create rsvp-ingress-unidir <rsvp-ingress-unidir> dest-ip <IP Address> backup-tunnel <static-ingress-corout></pre> where backup-tunnel is the static ingress co-routed tunnel created in step 1 <static-ingress-corout>
3	Create a static associated bidirectional tunnel with the dynamic ingress unidirectional tunnel configured as the forward tunnel: <pre>gmpls tp-tunnel create static-ingress-assoc <static-ingress-assoc> forward-tunnel <rsvp-ingress-unidir> reverse-dyntun-name <reverse-dyntun-name> bfd-monitor <enable disable></pre> where forward-tunnel is the dynamic ingress unidirectional tunnel created in step 2. <rsvp-ingress-unidir>

- 4 Create an LDP-signaled virtual circuit with peer set to the destination address of the static ingress co-routed tunnel and the ingress transport co-routed primary TP tunnel set to the static ingress co-routed tunnel:

```
mpls l2-vpn create dynamic-vc <dynamic-vc> pw-id <NUMBER:
1-2147483647> peer <IP Address> tp-tunnel-ingr-corout
<static-ingress-corout> pw-cword on pw-mode <mesh |
spoke>]
```

where

tp-tunnel-ingr- is the static ingress co-routed tunnel created in [step 1](#)
corout <static-
ingress-corout>

- 5 Attach the virtual circuit to a virtual switch:

```
virtual-switch attach mpls-vc <dynamic-vc> vs <vs>
—end—
```

Example

The following commands configure a static ingress co-routed tunnel named tp3-in-co-1 as the primary tunnel and a dynamic ingress unidirectional tunnel named tp3-rsvp-1 as the backup tunnel.

```
gmpls tp-tunnel create static-ingress-corout tp3-in-co-1 dest-ip 10.10.10.10
next-hop-ip 166.166.166.100 forward-out-label 251 reverse-in-label 252 bfd-
monitor enable
gmpls tp-tunnel create rsvp-ingress-unidir tp3-rsvp-1 dest-ip 10.10.10.10
backup-tunnel tp3-in-co-1
gmpls tp-tunnel create static-ingress-assoc tp-assoc-1 forward-tunnel tp3-
rsvp-1 reverse-dyntun-name tp1-rsvp-1 bfd-monitor enable
mpls l2-vpn create dynamic-vc dvc1-1 pw-id 300 peer 10.10.10.10 tp-tunnel-
ingr-corout tp3-in-co-1 pw-cword on pw-mode spoke
virtual-switch ethernet create vs vs-data-1
virtual-switch ethernet attach mpls-vc dvc1-1 vs vs-data-1
```

Procedure 10-5

Configuring a dynamic associated tunnel as primary and a static co-routed tunnel as backup

Configure a dynamic associated tunnel as primary with a static co-routed tunnel as backup if required for the network configuration.

Step	Action
1	Create a dynamic associated tunnel: <pre>gmpls tp-tunnel create rsvp-ingress-unidir <rsvp-ingress-unidir> dest-ip <IP Address></pre> where rsvp-ingress-unidir <rsvp-ingress-unidir> is the name of the dynamic ingress unidirectional tunnel dest-ip <IP Address> is the destination IP address
2	Create a static ingress co-routed tunnel: <pre>gmpls tp-tunnel create static-egress-corout <static-egress-corout> src-ip <IP Address> prev-hop-ip <IP Address> forward-in-label <NUMBER: 16-1044479> reverse-out-label <NUMBER: 16-1044479> backup-tunnel <MPLS static egress primary tp-tunnel></pre> where src-ip <IP Address> is the IP address specified as the destination address for the dynamic associated tunnel created in step 1.
3	Create a static ingress associated tunnel with the dynamic associated tunnel as the forward tunnel: <pre>gmpls tp-tunnel create static-ingress-assoc <static-ingress-assoc> forward-tunnel <rsvp-ingress-unidir> reverse-dyntun-name <reverse-dyntun-name> bfd-monitor <enable disable></pre> where forward-tunnel <rsvp-ingress-unidir> is the dynamic associated tunnel created in step 1.

- 4 Create an LDP-signaled virtual circuit with peer set to the destination address of the dynamic associated tunnel and the static egress transport co-routed primary TP-Tunnel set to the static ingress associated tunnel:

```
mpls l2-vpn create dynamic-vc <dynamic-vc> pw-id <NUMBER:
1-2147483647> tp-tunnel-assoc <static-ingress-assoc>
peer <IP Address>
```

where

tp-tunnel-assoc is the static ingress associated tunnel created in [step 3](#)
<static-ingress-
assoc>

- 5 Attach the virtual circuit to a virtual switch:

```
virtual-switch attach mpls-vc <dynamic-vc> vs <vs>
```

—end—

Example

The following commands configure a dynamic associated tunnel named tp1-rsvp-3 as the primary tunnel and a static egress co-routed tunnel named tp3-eg-co-3 as the backup tunnel.

```
gmpls tp-tunnel create rsvp-ingress-unidir tp1-rsvp-3 dest-ip 40.40.40.40
gmpls tp-tunnel create static-egress-corout tp3-eg-co-3 src-ip 40.40.40.40
prev-hop-ip 166.166.166.200 forward-in-label 271 reverse-out-label 272
backup-tunnel tp1-rsvp-3
gmpls tp-tunnel create static-ingress-assoc tp-assoc-3 forward-tunnel tp1-
rsvp-3 reverse-dyntun-name tp-rsvp-3 bfd-monitor enable
mpls l2-vpn create dynamic-vc dvc-3 pw-id 151 tp-tunnel-assoc tp-assoc-3 peer
40.40.40.40
virtual-switch ethernet create vs vs-data-1 mode vpls
virtual-switch ethernet attach mpls-vc dvc-3 vs vs-data-1
```

Procedure 10-6

Configuring a primary tunnel with an explicit path and a backup tunnel without an explicit path

Configure a primary dynamic ingress unidirectional tunnel with an explicit path and a backup tunnel without an explicit path as indicated by network requirements.

Step	Action
1	Create a dynamic ingress unidirectional tunnel with an explicit path: <pre>gmpls tp-tunnel create rsvp-ingress-unidir <rsvp-ingress-unidir> dest-ip <IP Address> explicit-tunnel-path <MPLS Rsvp Path></pre> where rsvp-ingress-unidir <rsvp-ingress-unidir> is the name of the dynamic ingress unidirectional tunnel dest-ip <IP Address> is the destination IP address explicit-tunnel-path <MPLS Rsvp Path> is the explicit path name for the tunnel
2	Create the backup tunnel for the dynamic ingress unidirectional tunnel: <pre>gmpls tp-tunnel create rsvp-ingress-unidir <rsvp-ingress-unidir> dest-ip <IP Address> backup-tunnel <MPLS ingress tp-unidirtunnel></pre> where rsvp-ingress-unidir <rsvp-ingress-unidir> is the name of the backup dynamic ingress unidirectional tunnel dest-ip <IP Address> is the same destination address as for the primary dynamic ingress unidirectional tunnel created in step 1. backup-tunnel <MPLS ingress tp-unidirtunnel> specify the dynamic ingress unidirectional tunnel created in step 1.

—end—

Example

The following example configures a primary dynamic ingress unidirectional tunnel named `rsvp_test` with an explicit path of `path1`. The backup tunnel is named `rsvp_test_bkp`.

```
gmpls tp-tunnel create rsvp-ingress-unidir rsvp_test dest-ip  
21.21.21.21 explicit-tunnel-path path1
```

```
gmpls tp-tunnel create rsvp-ingress-unidir rsvp_test_bkp dest-  
ip 21.21.21.21 backup-tunnel rsvp_test
```

Procedure 10-7

Configuring a primary tunnel without an explicit path and a backup tunnel with an explicit path

Configure a primary dynamic ingress unidirectional tunnel without an explicit path and a backup dynamic ingress unidirectional tunnel with an explicit path as indicated by network requirements.

Step	Action
1	Create a dynamic ingress unidirectional tunnel: <pre>gmpls tp-tunnel create rsvp-ingress-unidir <rsvp-ingress-unidir> dest-ip <IP Address></pre> where rsvp-ingress-unidir <rsvp-ingress-unidir> is the name of the dynamic ingress unidirectional tunnel dest-ip <IP Address> is the destination IP address
2	Create a backup tunnel with explicit path for the dynamic ingress unidirectional tunnel: <pre>gmpls tp-tunnel create rsvp-ingress-unidir <rsvp-ingress-unidir> dest-ip <IP Address> backup-tunnel <MPLS ingress tp-unidirtunnel> explicit-tunnel-path <MPLS Rsvp Path></pre> where rsvp-ingress-unidir <rsvp-ingress-unidir> is the name of the backup dynamic ingress unidirectional tunnel dest-ip <IP Address> is the same destination address as for the primary dynamic ingress unidirectional tunnel created in step 1. backup-tunnel <MPLS ingress tp-unidirtunnel> specify the dynamic ingress unidirectional tunnel created in step 1. explicit-tunnel-path <MPLS Rsvp Path> is the explicit path name for the tunnel

—end—

Example

The following example configures a primary dynamic ingress unidirectional tunnel named `rsvp_test`. The backup tunnel is named `rsvp_test_bkp` and has an explicit path of `path2`.

```
gmpls tp-tunnel create rsvp-ingress-unidir rsvp_test dest-ip  
21.21.21.21
```

```
gmpls tp-tunnel create rsvp-ingress-unidir rsvp_test_bkp dest-  
ip 21.21.21.21 backup-tunnel rsvp_test explicit-tunnel-path  
path2
```

Procedure 10-8

Configuring a primary tunnel and a backup tunnel without explicit paths

Configure a primary dynamic ingress unidirectional tunnel and a backup dynamic ingress unidirectional tunnel without explicit paths as indicated by network requirements.

Step	Action
1	Create a dynamic ingress unidirectional tunnel: <pre>gmpls tp-tunnel create rsvp-ingress-unidir <rsvp-ingress-unidir> dest-ip <IP Address></pre> where rsvp-ingress-unidir <rsvp-ingress-unidir> is the name of the dynamic ingress unidirectional tunnel dest-ip <IP Address> is the destination IP address
2	Create a backup tunnel for the dynamic ingress unidirectional tunnel: <pre>gmpls tp-tunnel create rsvp-ingress-unidir <rsvp-ingress-unidir> dest-ip <IP Address> backup-tunnel <MPLS ingress tp-unidirtunnel></pre> where rsvp-ingress-unidir <rsvp-ingress-unidir> is the name of the backup dynamic ingress unidirectional tunnel dest-ip <IP Address> is the same destination address as for the primary dynamic ingress unidirectional tunnel created in step 1. backup-tunnel <MPLS ingress tp-unidirtunnel> specify the dynamic ingress unidirectional tunnel created in step 1.

—end—

Example

The following example configures a primary dynamic ingress unidirectional tunnel named rsvp_test and a backup tunnel named rsvp_test_bkp.

```
gmpls tp-tunnel create rsvp-ingress-unidir rsvp_test dest-ip 21.21.21.21
```

10-24 Additional tunnel configuration

```
gmpls tp-tunnel create rsvp-ingress-unidir rsvp_test_bkp dest-  
ip 21.21.21.21 backup-tunnel rsvp_test
```

LDP configuration

LDP is required for dynamic MPLS deployments for signaling virtual circuits. Sessions are established using UDP port 646 to a peer IP address that can be directly connected or several hops away.

By default, LDP is globally disabled. For dynamic deployments, it must be globally enabled.

LDP supports authentication to prevent unwanted connections.

This chapter provides the following procedures:

- [“Configuring LDP globally”](#)
- [“Displaying LDP configuration”](#)
- [“Configuring LDP authentication”](#)
- [“LDP LSP configuration”](#)
- [“Creating an MPLS TP-TE gateway”](#)

Procedure 11-1

Configuring LDP globally

You can

- enable LDP globally
- disable LDP globally
- modify global LDP attributes

Step	Action

To enable LDP globally

1	Enable LDP globally: <code>ldp enable</code>

To disable LDP globally

2	Disable LDP globally: <code>ldp disable</code>

To modify global LDP attributes

3	Modify global LDP attributes: <code>ldp set {[hello-hold-time <NUMBER: 1-65535>], [keep- alive-hold-time <NUMBER: 1-65535>]}</code> <code>—end—</code>

Procedure 11-2

Displaying LDP configuration

Display LDP configuration.

Step	Action

1	Display LDP configuration: <pre>ldp show [adjacency] [sessions] [timers]</pre> —end—

Example

The following example shows a global configuration.

```
> ldp show
```

LDP GLOBAL CONFIG	
Parameter	Value
LDP Admin State	Disabled
LDP Oper State	Disabled

Procedure 11-3

Configuring LDP authentication

You can

- add an IP entry and password
- remove an IP entry and password
- display IP entries with an encoded password

Step	Action

To add an IP entry and password

- 1 Add an IP entry and password:

```
ldp authentication set {peer <IPAddress>}
{authentication-type <md5>} {password
<PasswordString[31]>} {secret <SecretString[62]>}
```

To remove an IP entry and password

- 2 Remove an IP entry and password:

```
ldp authentication unset {peer <IPAddress>}
```

To display IP entries with an encoded password

- 3 Display IP entries with an encoded password:

```
ldp authentication show
—end—
```

Example

The following example adds an IP entry and password.

```
ldp authentication set peer 1.2.3.4 password myPassword
ldp authentication show
+----- LDP Authentication Configuration Summary -----+
| Router Id | Encoded Password |
+-----+-----+
| 1.2.3.4 | ffe5109e033a3716e94f |
+-----+-----+
```

Procedure 11-4

LDP LSP configuration

For LDP LSP establishment, additional to the global LDP enable command, LDP needs to be enabled on the interfaces between the directly connected LDP neighbors. This enables the nodes to form LDP adjacency and session for exchange of labels for the LDP LSP. LDP runs on top of IGP such as OSPF/IS-IS, so LDP and the IGP must be configured/enabled on the same set of interfaces.

Requirements

LDP must be enabled globally.

- 1 Use this command to enable LDP on an interface for LDP LSP establishment:
`ldp lsp enable ip-interface <ip_interface_name>`
- 2 Use this command to set the hello hold time value for Topology LDP. Default hello-hold-time is 15 seconds:
`ldp lsp set hello-hold-time <SECONDS: 15..65535>`
- 3 Use this command to set label-type in global ldp database. Default value is implicit-null.
`ldp set label-type <implicit-null | non-reserved>`
- 4 Use this command to set the keep alive hold time value for the Topology LDP. Default keep-alive-hold-time is 40 seconds:
`ldp lsp set keep-alive-hold-time <SECONDS: 40..65535>`
- 5 Use this command to display the MPLS LDP LSP table:
`mpls ldp-lsp show`
- 6 Use this command to show the LDP LSP timers:
`ldp lsp show timers`
- 7 Use this command to show LDP LSP in-label table:
`ldp lsp show in-label`
- 8 Use this command to show LDP LSP out-label table:
`ldp lsp show out-label`

Example

The following example enables LDP on an interface for LDP LSP establishment and displays MPLS data.

```
> ldp lsp enable ip-interface ip31
> mpls ldp-lsp show
```

MPLS LDP LSP TABLE							
Type	Tunnel Index	VIF Index	Tunnel Name / FEC	In Label	Out Label	Next Hop	Status
Ingress	1	262145	192.168.200.5/32	-	3	192.168.35.2	ENA

11-6 LDP configuration

```
+-----+-----+-----+-----+-----+-----+-----+-----+
> ldp lsp show timers
+-----+-----+-----+-----+-----+-----+-----+-----+
| Parameter                                | Value |
+-----+-----+-----+-----+-----+-----+-----+-----+
| LDP Hello Hold Time (sec)               | 15    |
| LDP Keepalive Hold Time (sec)           | 40    |
+-----+-----+-----+-----+-----+-----+-----+-----+

> ldp lsp show in-label
+-----+-----+-----+-----+-----+-----+-----+-----+
| LDP IPv4 FECs                           | Peer ID | In Label | State |
|                                           | (to)    | (Sent)  |      |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 192.168.200.3/32                       | 192.168.200.5:0 | 3       | Active |
+-----+-----+-----+-----+-----+-----+-----+-----+

> ldp lsp show out-label
+-----+-----+-----+-----+-----+-----+-----+-----+
| LDP IPv4 FECs                           | Peer ID | Out Label | State |
| Next Hop IP                           | (from)  | (Rcvd)   |      |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 192.168.200.5/32                       | 192.168.200.5:0 | 3       | Active |
| 192.168.35.2                           |                  |         |      |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

Procedure 11-5

Creating an MPLS TP-TE gateway

Create an MPLS TP-TE gateway to connect the TP and TE MPLS domains in order to provide end-to-end service resiliency.

Gateway mode is set when the virtual circuit is created. You cannot modify the PW mode of a virtual circuit set to gateway until the virtual circuit is removed from the virtual switch.

The virtual switch that the two gateway virtual circuits attach to cannot have other virtual circuits attached to it.

Only unprotected virtual circuits can be used in gateway mode. The status TLV attribute must be set to on.

Step	Action

1	Create the gateway PWE:

```
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> tp-tunnel-egrs-
corout <MPLS egress primary tp corout tunnel> [pw-mode
<gateway | mesh | spoke |switching>] [status-tlv <on |
off>]
```

Example

The following command creates static PWE gateway StGw1.

```
> mpls l2-vpn create static-vc S1S3_95 pw-id 378 peer 192.168.xxx.1 in-label
65726 out-label 65726 tp-tunnel-egrs-corout-static S1S3 pw-mode gateway
status-tlv on
> mpls l2-vpn create dynamic-vc S3S5_95 pw-id 380 peer 192.168.xxx.5 ldp-lsp
pw-mode gateway status-tlv on
> virtual-switch interface attach mpls-vc S3S5_95 vs S1S3S5_95
> virtual-switch interface attach mpls-vc S1S3_95 vs S1S3S5_95
> mpls l2-vpn create static-vc S1S3_96 pw-id 382 peer 192.168.xxx.1 in-label
65728 out-label 65728 tp-tunnel-egrs-corout-static S1S3 pw-mode gateway
status-tlv on
```

The following command displays gateway StGw1.

```
> mpls l2-vpn show vc S1S3_96
```

Pseudowire ID	378
Customer Name	
Pseudowire Name	S1S3_96
Signaling Type	Static

11-8 LDP configuration

Pseudowire Admin State	Enabled
Operational State	Up
Pseudowire Blocking	Yes
Remote Peer IP Address	192.168.xxx.1
Incoming Label	65728
Outgoing Label	65728
Status TLV	On
Status Query	On
Refresh Status Interval	600 seconds
Local Fault	None
Remote Fault	PW NFwd
MTU Size	1500
Pseudowire Type	Ethernet Raw
Pseudowire Mode	Gateway
Forward CoS Profile Name	DefaultVcCosProfile
Egress L2PT Transform	Disabled
Forward Vccv Profile Name	DefaultPwVccvProfile
Local CC:CV	n/a
Remote CC:CV	n/a
Operating CC:CV	Gal-GAch : LSP
Tunnel Virtual Interface	S1S3S5_95
Virtual Switch Name	S1S3S5_488
Failure Reason	n/a
Multi-Chassis	Off

OAM Signalling Statistics	

Status Frames Sent	2081
Status Frames Received	2068
Status Acks Sent	2067
Status Acks Received	2081
Status Bad PDUs Received	0
Status Bad Acks Received	0
Mac Withdrawal Frames Sent	0
Mac Withdrawal Frames Received	16
Mac Withdrawal Acks Sent	16
Mac Withdrawal Acks Received	0
Mac Withdrawal Bad PDUs Received	0
Mac Withdrawal Bad Acks Received	0
Mac Withdrawal Sent Sequence Number	0
Mac Withdrawal Received Sequence Number	15

Virtual circuit configuration

When configuring primary virtual circuits, follow these configuration rules:

- All virtual circuit names must be unique.
- A virtual circuit cannot be deleted unless it is detached from the virtual switch.
- Dynamic virtual circuits: For a given peer node, the Pseudowire ID of the virtual circuits must be unique.
- Static virtual circuits: Ingress label is unique across virtual circuits and tunnels.

Note: MPLS virtual circuits share internal resources with MPLS attachment circuits and sub-ports. This means that the scale of these features may be affected by the number of MPLS virtual circuit and attachment circuit instances on the switch.

Primary dynamic and static MPLS virtual circuits support the attributes listed in [Table 12-1](#).

Table 12-1
Primary virtual circuit attributes

Attribute	Description
Name	Unique 31 character name.
PW-ID	Pseudowire identifier. Ranging from 1-4294967295. This identifier must be the same at both ends of the PW.
FEC 129	Also known as Generalized ID FEC. FEC 129 is advertised in Targeted LDP to exchange the PW label for the PW FEC. This is the preferred mechanism for exchanging PW FEC.
Peer	Peer IP address.
Ingress label	Ingress label for the virtual circuit. Only applies to static virtual circuits.
Egress label	Egress label for the virtual circuit. Only applies to static virtual circuits.
Tunnel	Name of the tunnel.

12-2 Virtual circuit configuration

Attribute	Description
PW type	Pseudowire type. For static virtual circuits, the types are raw or tagged. Default is raw. For dynamic virtual circuits, the types are eth-raw, eth-tagged, or tdm.
PW mode	Pseudowire mode: • Mesh. A virtual circuit between two PEs. • Spoke. A virtual circuit emanating towards the MTU-s. • Switching: Configures multisegment pseudowire (MS-PW).
PW reversion	Pseudowire reversion. When pseudowire reversion is on, the system switches from the backup to the primary tunnel after the fault on the primary is cleared.
Pseudowire CoS Profile	Pseudowire CoS profile name.
Pseudowire VCCV Profile	Pseudowire VCCV profile name.
MTU	Maximum transmission unit 1500-9128. Default is 1500. This parameter is exchanged during dynamic establishment of a PW and must match in order for the PW to become operational between two PEs. Note: Mismatch of the MTU is one of the most common reasons for PW to fail to become operational.
Status TLV	On or off. Applies to virtual circuits and static virtual circuits. The static virtual circuit carries status TLV in the OAM channel.

For dual-homed protection of an MTU-s primary dynamic virtual circuit, the secondary virtual circuit is defined by setting the peer IP address, tunnel, primary PW-ID, and secondary PW-ID. All other attributes, other than control word, are inherited from the primary.

For dual-homed protection, it is important to ensure that Pseudowire Status messaging is enabled for both the primary and the secondary pseudowires. This is done by setting the status-tlv attribute to on at both endpoints of the primary virtual circuits, as well as on the remote PE node of the secondary virtual circuit.

By default, the primary virtual circuit is the active virtual circuit. If the primary fails, the secondary virtual circuit becomes the active virtual circuit. The active virtual circuit (whether it is the primary or secondary) remains active unless it fails or it can be manually switched over.

Changing the PW control word

When the pw-cword setting needs to be changed, the PW must be administratively disabled on both endpoints first, followed by the set operation described in [“Modifying dynamic virtual circuits” on page 12-9](#).

PW group, infrastructure PW, and PW group switchover support

The new PW group feature allows pseudowires to be grouped together by assigning a “pw-group-id”. The PW group is created automatically when any first PW is created with “pw-group-id” set. There is no separate CLI to create PW group. Primary PWs which share the same tunnel can be configured to be part of the same PW-group. Protection PWs inherit the pw-group-id from their primary PW and can have their tunnels different from the other protected PW members of the group.

Infrastructure PW (also referred as Infra PW) is a special purpose PW introduced in conjunction with PW-group and is used to carry OAM traffic only (such as VCCV BFD or PW Status Tlv) and not user data. The Infra PW is a member of a PW-group and the decision for the PW-group switchover is governed by its state. Only one Infra PW is allowed per PW-group. Configuration of VCCV-BFD and Status-Tlv is supported on Infra PW. (For examples showing how to configure VCCV BFD on a PW, see [VCCV BFD over static SS-PW over single LSP example](#) and [VCCV BFD over static MS-PW over single LSP example](#).) The configuration of a backup PW for the infrastructure PW is not permitted. The configuration of VCCV BFD over non-infra members of the group is not permitted.

The pseudowire group switchover is a mechanism triggered by the Infra PW fault. It allows all primary pseudowires in a group to be quickly switched over to backup pseudowires. It reduces the switchover time from primary pseudowires to backup pseudowires when a fault is encountered and thus results in less service disruption. The reduced switchover time is achieved by sending a single signal to the datapath to switch all the primary PWs of a group at once. During the PW group switchover, MAC withdrawal is sent on all the backup PWs of the PW group.

In a PW group, there is a slight change in behavior of status signaling for backup PWs. Prior to the concept of the PW group, the backup PWs always sent NFW status to their remote ends. With PW groups, the backup PWs of a group now send FWD status to the remote end, if there is no fault on the backup link. Now the remote PW segments (at SPEs) of the backup PWs of the group (at MTUs) are in a FWD state to accept the switched-over traffic (during group switchover at MTUs).

Infrastructure PW and PW group switchover is supported for both SS-PW and MS-PW.

Note: For MS-PW, it is the operator's responsibility to configure PW segments on the required LSP of the SPE nodes. All the PW segments whose remote PWs belong to a particular group must be configured on the same LSP at SPE.

In case of fault clearance on the primary link, WTR (wait to restore) timers of the non-Infra PWs are kicked off only when the Infra PW fault is cleared. Reversion of all PWs in a group occurs individually. Group level reversion is not supported.

The PW group, infrastructure PW and PW group switchover features are supported only for static PW.

Logging for PW state change/up time

Logging of the MPLS pseudowire (PW) switchover state change is now available. It can be traced in NMS alarms or the NE syslog. These logs and the PW up time help in debugging/diagnosis of PW network issues.

As a result, two new SNMP traps are defined:

- The first trap is raised when the PW changes from a blocking to a forwarding state.
- The second trap is raised when the PW changes from a forwarding to a blocking state.

Note: These SNMP traps are not raised for switching PWs.

The time for which the pseudowire remains in a forwarding state during switchover is shown in the detailed output show of "mpls l2-vpn show vc". (It can also be fetched through an SNMP get operation.) The new Forwarding State Up Time parameter is visible only in the case of forwarding virtual circuits.

Virtual circuit statistics at the S-PE NE node

Traffic statistics for S-PE NE VCs are now supported. The egress traffic statistics for switching VCs are fetched from the egress object created in the device and the Ingress traffic statistics for switching VCs are fetched from the incoming VC label. The monitoring of the VC traffic statistics is available after both segment S-PE NE VCs are stitched in the device and operationally up. The traffic statistics work for VCs provisioned over both protected and unprotected tunnels.

The traffic statistics can be requested on demand. MPLS OAM statistics are accounted for S-PE NE node VCs. Statistics for S-PE node VCs are displayed in CLI in a similar way to existing LER VCs.

When the VC present in one segment becomes operationally down or detached from the VS, it impacts another segment as follows:

- Tx and Rx traffic statistics are reset to 0 for both segment VCs.
- No increment of Rx or Tx statistics is done for the other segment.

This chapter provides the following procedures:

- [“Configuring dynamic virtual circuits” on page 12-6](#)
- [“Modifying dynamic virtual circuits” on page 12-9](#)
- [“Configuring static virtual circuits” on page 12-10](#)
- [“Modifying static virtual circuits” on page 12-14](#)
- [“Displaying virtual circuits” on page 12-15](#)
- [“Configuring L2-VPN virtual circuit CoS profile” on page 12-23](#)
- [“Displaying L2-VPN virtual circuit CoS profiles” on page 12-24](#)
- [“Configuring virtual circuit VCCV profiles” on page 12-25](#)
- [“Displaying virtual circuit VCCV profiles” on page 12-28](#)
- [“Configuring MS-PW stitching” on page 12-29](#)
- [“Configuring PW-group protection example” on page 12-32](#)

Procedure 12-1

Configuring dynamic virtual circuits

You can

- create a dynamic virtual circuit
- create a dynamic protection virtual circuit

Step	Action

To create a dynamic virtual circuit

- 1 Create a dynamic virtual circuit:

```
mpls l2-vpn create dynamic-vc <dynamic-vc> {{pw-id
<NUMBER: 1... 2147483647> | fec-129 {[agi <String>] [saii
<String>] [taii <String>]]}} {peer <IP Address>} {te-
tunnel <MPLS ingress primary tunnel>} {tp-tunnel-ingr-
corout <MPLS ingress primary tp corout tunnel>} {tp-
tunnel-egress-corout-static <MPLS static egress primary
tp-tunnel>} {tp-tunnel-egrs-corout-dynamic <String>}
{te-tunnel-assoc <MPLS assoc te-tunnel>} {tp-tunnel-assoc
<MPLS assoc tp-tunnel>} [pw-cword <on|off>] [pw-type <eth-
raw|eth-tagged|tdm>] [mtu <1500-9128>] [status-tlv
<on|off>] [pw-mode <mesh | spoke|switching>] [pw-
reversion <on |off>] [reversion-hold-time <SECONDS:
0..3600>] [pw-cos-profile <MPLS Pseudowire COS Profile>]
[pw-vccv-profile <MPLS Pseudowire VCCV Profile>]
```

where

dynamic-vc is the name of the dynamic virtual circuit.
<dynamic-vc>

pw-id<NUMBER: is the VPN identifier.
1-2147483647>

fec-129 selects FEC 129, which facilitates single-sided
configuration.

agi <String> is the Attachment Group Identifier, which specifies the VPN
identifier string.

saii <String> is the Source Attachment Individual Identifier, which
specifies the local attachment identifier string.

taii <String> is the Target Attachment Individual Identifier, which specifies
the remote attachment identifier string.

peer <IP
Address> is the destination IP address.

where

te-tunnel <MPLS ingress primary tunnel>	is the ingress transport primary TE tunnel.
tp-tunnel-ingr-corout <MPLS ingress primary tp corout tunnel>	is the ingress transport co-routed primary TP tunnel.
tp-tunnel-egrs-corout-static <MPLS static egress primary tp-tunnel>	is the static egress transport co-routed primary TP tunnel.
tp-tunnel-egrs-corout-dynamic <String>	is the name of the dynamic egress transport co-routed primary TP tunnel.
te-tunnel-assoc <MPLS assoc te-tunnel>	is the ingress transport associated primary TE tunnel.
tp-tunnel-assoc <MPLS assoc tp-tunnel>	is the ingress transport associated primary TP tunnel.
pw-cword <on off>	enables or disables pw control word. Default is off.
pw-type <eth-raw eth-tagged tdm>	is the pseudowire type.
mtu <NUMBER: 1500-9128>	is the MTU in bytes.
status-tlv <on off>	determines whether status TLV is on or off.
pw-mode <mesh spokel switching>	is the pseudowire mode.
pw-reversion <on off>	turns pw-reversion on or off. Default is on.

where

reversion-hold-time <SECONDS: 0..3600> sets the reversion-hold-time. Default is 30 sec. If the value is set to 0, there is no delay in switching over to the primary PW.

pw-cos-profile <MPLS Pseudowire CoS Profile> is the pseudowire COS profile name.

pw-vccv-profile <MPLS Pseudowire VCCV Profile> is the pseudowire VCCV profile name.

2 Enable the virtual circuit:

```
mpls l2-vpn enable vc <vc>
```

where

vc <vc> is the virtual circuit to enable.

To create a dynamic protection virtual circuit

3 Create a dynamic protection virtual circuit:

```
mpls l2-vpn protection create dynamic-vc <dynamic-vc>
{{secondary-pw-id <secondary-pw-id> | fec-129 {[agi
<String>] [saii <String>] [taii <String>]]}} {primary-vc-
name <Virtual Circuit Dynamic MPLS Name>} {peer <IP
address>} {te-tunnel <MPLS ingress primary tunnel>|tp-
tunnel-ingr-corout <MPLS ingress primary tp corout
tunnel>|tp-tunnel-egrs-corout-static <MPLS static egress
primary tp-tunne>|tp-tunnel-egrs-corout-dynamic
<String>} {te-tunnel-assoc <MPLS assoc te-tunnel>|tp-
tunnel-assoc <MPLS assoc tp-tunnel>}{pw-cword <on|off>}
```

—end—

Procedure 12-2

Modifying dynamic virtual circuits

Modify dynamic virtual circuits if network requirements change.

Step	Action

1	Modify dynamic virtual-circuit attributes

```
mpls l2-vpn set dynamic-vc <dynamic-vc> [pw-mode <mesh |
spoke>] [pw-cword <on|off>] [te-tunnel <MPLS ingress
primary tunnel>] [tp-tunnel-ingr-corout <MPLS ingress
primary tp corout tunnel>] [tp-tunnel-egrs-corout-static
<MPLS static egress primary tp-tunnel>] [te-tunnel-assoc
<MPLS assoc te-tunnel>] [tp-tunnel-assoc <MPLS assoc tp-
tunnel>] [customer <String[15]>] [pw-reversion <on |
off>] [reversion-hold-time <NUMBER: 0-3600>] [pw-vccv-
profile <MPLS Pseudowire VCCV Profile>] [pw-cos-profile
<MPLS Pseudowire CoS Profile>]
```

—end—

Procedure 12-3

Configuring static virtual circuits

Configure static virtual circuits.

Step	Action														
1	Create a static virtual circuit: <pre>mpls l2-vpn create static-vc <static-vc> {pw-id <NUMBER: 1-2147483647> peer <ip-addr-str>} {te-tunnel <Ingress-Primary-TE-Tunnel-Object> tp-tunnel-ingr-corout <Ingress-Primary-Corouted-TP-Tunnel-Object> tp-tunnel-egrs-corout-static <Static-Egress-Primary-Corouted-TP-Tunnel> tp-tunnel-egrs-corout-dynamic <NAME-STRING> te-tunnel-assoc <Primary-Associated-TE-Tunnel> tp-tunnel-assoc <Primary-Associated-TP-Tunnel>} {peer <IP address>} in-label <NUMBER: 16..1044479> out-label <NUMBER: 16..1048575> [pw-cword <on off>] [tunnel <lsp-name>] [pw-type <eth-raw eth-tagged tdm>] [tdm-profile <xml-tdm-profile>] [pw-cword <on off>] [status-tlv <on off>] [refresh-status-interval <0..65535>] [pw-mode <mesh spoke switching>] [pw-reversion <on off>] [reversion-hold-time <SECONDS: 0..3600>] [mtu <NUMBER: 1500-9128>] [pw-cos-profile <xml-mpls-pw-cos-profile>] [pw-vccv-profile <MPLS Pseudowire VCCV Profile>] [pw-group-id <group-id>] [is-infra-pw <yes no>]</pre> where <table> <tr> <td>static-vc <static-vc></td><td>Static virtual circuit name</td></tr> <tr> <td>{pw-id <NUMBER: 1-2147483647>}</td><td>Set VPN ID</td></tr> <tr> <td>{peer <IP address>}</td><td>Set Destination IP address</td></tr> <tr> <td>{te-tunnel <Ingress-Primary-TE-Tunnel-Object></td><td>Selects ingress transport primary TE-Tunnel.</td></tr> <tr> <td>tp-tunnel-ingr-corout <Ingress-Primary-Corouted-TP-Tunnel-Object></td><td>Selects ingress transport co-routed primary TP-Tunnel.</td></tr> <tr> <td>tp-tunnel-egrs-corout-static <Static-Egress-Primary-Corouted-TP-Tunnel></td><td>Selects static egress transport co-routed primary TP-Tunnel.</td></tr> <tr> <td>tp-tunnel-egrs-corout-dynamic <NAME-STRING></td><td>selects dynamic egress transport co-routed primary TP-Tunnel name.</td></tr> </table>	static-vc <static-vc>	Static virtual circuit name	{pw-id <NUMBER: 1-2147483647>}	Set VPN ID	{peer <IP address>}	Set Destination IP address	{te-tunnel <Ingress-Primary-TE-Tunnel-Object>	Selects ingress transport primary TE-Tunnel.	tp-tunnel-ingr-corout <Ingress-Primary-Corouted-TP-Tunnel-Object>	Selects ingress transport co-routed primary TP-Tunnel.	tp-tunnel-egrs-corout-static <Static-Egress-Primary-Corouted-TP-Tunnel>	Selects static egress transport co-routed primary TP-Tunnel.	tp-tunnel-egrs-corout-dynamic <NAME-STRING>	selects dynamic egress transport co-routed primary TP-Tunnel name.
static-vc <static-vc>	Static virtual circuit name														
{pw-id <NUMBER: 1-2147483647>}	Set VPN ID														
{peer <IP address>}	Set Destination IP address														
{te-tunnel <Ingress-Primary-TE-Tunnel-Object>	Selects ingress transport primary TE-Tunnel.														
tp-tunnel-ingr-corout <Ingress-Primary-Corouted-TP-Tunnel-Object>	Selects ingress transport co-routed primary TP-Tunnel.														
tp-tunnel-egrs-corout-static <Static-Egress-Primary-Corouted-TP-Tunnel>	Selects static egress transport co-routed primary TP-Tunnel.														
tp-tunnel-egrs-corout-dynamic <NAME-STRING>	selects dynamic egress transport co-routed primary TP-Tunnel name.														

where

te-tunnel-assoc <Primary-Associated-TE-Tunnel>	selects ingress transport associated primary TE-Tunnel.
tp-tunnel-assoc <Primary-Associated-TP-Tunnel>	selects ingress transport associated primary TP-Tunnel.
{ingress-label <NUMBER: 16-1044479>}	sets the MPLS decap label.
{egress-label <NUMBER: 16-1048575>}	sets the MPLS encap label.
pw-cword <on off>	enables or disables pw control word. Default is off.
{tunnel <MPLS ingress primary tunnel>}	transport tunnel
[pw-type <eth-raw eth-tagged tdm>]	selects the pseudowire type. pw-type must be the same for static to dynamic PWs.
[tdm-profile <xml-tdm-profile>]	Sets TDM profile to be used to setup a static TDM pseudowire.
status-tlv <on off>	determines if status TLV is on or off.
status-interval <0..65535>	refreshes the current status of the pseudowire to ensure that each end has the other's correct pseudowire status. 0 indicates no status refresh. The default is 600 seconds. Note: The status interval value should consider how many PWs are expected to be configured. For a large number of PWs, the refresh interval must be set to a higher value. This reduces the frequency of a large number of PW status message exchanges.
[pw-mode <mesh spoke switching>]	selects the pseudowire mode.
pw-reversion <on off>	turns pw-reversion on or off. Default is on.
reversion-hold-time <SECONDS: 0..3600>	sets the reversion-hold-time. Default is 30 sec. If the value is set to 0, there is no delay in switching over to the primary PW.
[mtu <NUMBER: 1500-9128>]	sets MTU (bytes). The MTU must be the same for static to dynamic PWs
[pw-cos-profile <MPLS Pseudowire CoS Profile>	is the pseudowire VCCV profile name.

where

[pw-vccv-profile <MPLS Pseudowire VCCV Profile>] is the pseudowire VCCV profile name.

[pw-group-id <group-id>] is the pseudowire group identifier. Default is 0 (not part of a pseudowire group).

[is-infra-pw <yes | no>] specifies whether this is an infrastructure pseudowire. Default is no.

2 Enable the virtual circuit:

```
mpls l2-vpn enable [vc <vc-name>]
```

where

vc <vc-name> is the virtual circuit to enable.

3 Create a static protection virtual circuit:

```
mpls l2-vpn protection create static-vc <static-vc>
{secondarypw-id <NUMBER: 1-2147483647>} {primary-vc-name
<Virtual Circuit Static MPLS Name>} {peer <IP address>}
{in-label <NUMBER: 16-1044479>} {out-label <NUMBER: 16-
1044479>} {pw-cword <on|off>} {te-tunnel <MPLS ingress
primary tunnel>|tp-tunnel-ingr-corout <MPLS ingress
primary tp corout tunnel>|tp-tunnel-egrs-corout-static
<MPLS static egress primary tp-tunnel>|tp-tunnel-egrs-
corout-dynamic <String>|te-tunnel-assoc <MPLS assoc te-
tunnel>|tp-tunnel-assoc <MPLS assoc tp-tunnel>} [refresh-
status-interval <NUMBER: 0-65535>]
```

where

static-vc <static-vc> static virtual circuit name.

{secondarypw-id <NUMBER: 1-2147483647>} sets the VPN ID.

{peer <IP address>} sets the Destination IP address.

{in-label <NUMBER: 16-1044479>} sets MPLS decapsulation label.

{out-label <NUMBER: 16-1044479>} sets MPLS encapsulation label.

pw-cword <on|off> enables or disables pw control word. Default is off.

{te-tunnel <MPLS ingress primary tunnel>} selects the ingress transport primary TE-Tunnel.

{tp-tunnel-ingr-corout <PLS ingress primary tp corout tunnel>} selects the ingress transport co-routed primary TP-Tunnel.

{tp-tunnel-egrs-corout-static <MPLS static egress primary tp-tunnel>} selects the static egress transport co-routed primary TP-Tunnel.

where

{tp-tunnel-egrs-corout-dynamic <String>}	the dynamic egress transport co- routed TP-Tunnel name.
{te-tunnel-assoc <MPLS assoc te- tunnel>}	selects the ingress transport associated primary TE-Tunnel.
tp-tunnel-assoc <MPLS assoc tp- tunnel>}	selects the ingress transport associated primary TP-Tunnel.
refresh-status-interval <Number: 0-65535>	sets the refresh timer interval.

—end—

Examples

The following example configures multi-segment pseudowire.

```
mpls 12-vpn create static-vc SVCoS-21-00 pw-id 120000 peer 1.1.1.1.1 in-label  
12000 out-label 21000 pw-cword on tp-tunnel-ingr-corout stp-co-21-00 pw-type  
eth-tagged status-tlv on pw-mode switching
```


Procedure 12-4

Modifying static virtual circuits

Modify static virtual circuits if network requirements change.

Step	Action

1	Modify static virtual circuits:

```
mpls l2-vpn set static-vc <static-vc> [pw-mode <mesh |  
spoke>] [pw-cword <on|off>] [te-tunnel <MPLS ingress  
primary tunnel>] [tp-tunnel-ingr-corout <MPLS ingress  
primary tp corout tunnel>] [tp-tunnel-egrs-corout-static  
<MPLS static egress primary tp-tunnel>] [te-tunnel-assoc  
<MPLS assoc te-tunnel>] [tp-tunnel-assoc <MPLS assoc tp-  
tunnel>] [customer <String[15]>] [pw-reversion <on |  
off>] [reversion-hold-time <NUMBER: 0-3600>] [pw-vccv-  
profile <MPLS Pseudowire VCCV Profile>] [pw-cos-profile  
<MPLS Pseudowire CoS Profile>] [pw-group-id <group-id>]
```

—end—

Procedure 12-5

Displaying virtual circuits

You can display:

- all virtual circuits
- virtual circuits by attribute
- detailed output of a selected virtual circuit
- traffic statistics for a virtual circuit (incoming and outgoing packets, incoming and outgoing bytes), including multisegment pseudowire (MS-PW) and the segmented switching provider edge (S-PE) network element (NE) node virtual circuits
- pseudowires by customer name
- virtual circuit next hops
- virtual circuit information based on tunnel group
- virtual circuit resource summary
- virtual circuit groups configured with pw-group-id, including all the pseudowire members and Infra PW belonging to the pseudowire group

Step	Action

To display all virtual circuits

- 1 Display all Ethernet virtual circuits:

```
mpls l2-vpn show
```

To display virtual circuits by attribute

- 2 Display Ethernet virtual circuits by attribute:

```
mpls l2-vpn show matching-vc <matching-vc> {persist
<static|dynamic>} [state <up|down>] {pw-id <NUMBER>}
[destination <IP address>] [in-label <NUMBER: 16-
1044479>] [out-label <NUMBER: 16-1044479>] [recovery
<protected|unprotected>] [role
<primary|backup|standalone>] [te-tunnel <MPLS ingress
primary tunnel>] [tp-tunnel-ingr-corout <MPLS ingress
primary tp corout tunnel>] [tp-tunnel-egrs-corout-static
<MPLS static egress primary tp-tunnel>] {tp-tunnel-egrs-
corout-dynamic <Name (no!%,?*) String>} [te-tunnel-assoc
<MPLS assoc te-tunnel>] [tp-tunnel-assoc <MPLS assoc tp-
tunnel>] [pw-type <eth-raw|eth-tagged>] [pw-mode <mesh|
```

```
spoke|switching>] [mtu <#1500..9128>] [status-tlv
<off|on>] [pw-vccv-profile <MPLS Pseudowire VCCV
Profile>] [pw-cos-profile <MPLS Pseudowire COS Profile>]
```

where

matching-vc <matching-vc>	displays all matching VCs.
persist <static dynamic>	filters by persistence.
state <up down>	filters by operational state.
pw-id <NUMBER>	filters by VPN identifier.
destination <IP address>	filters by destination IP address of the virtual circuit.
in-label <NUMBER>	filters by inbound label value of the virtual circuit.
out-label <NUMBER>	filters by outbound label value of the virtual circuit.
recovery <protected unprotected>	filters by recovery type of the virtual circuit.
role <primary backup standalone>]	filters by current role of the virtual circuit.
te-tunnel <MPLS ingress primary tunnel>	displays ingress transport primary TE-Tunnels.
tp-tunnel-ingr- corout <MPLS ingress primary tp corout tunnel>	displays ingress transport co-routed primary TP-Tunnels.
tp-tunnel-egrs- corout-static <MPLS static egress primary tp-tunnel>	displays static egress transport corouted primary TP-Tunnels.
tp-tunnel-egrs- corout-dynamic <Name (no !%,?*) String>}	displays the dynamic egress transport co-routed primary TP tunnel.
[te-tunnel-assoc <MPLS assoc te- tunne>	displays ingress transport associated primary TE-Tunnels.

where

tp-tunnel-assoc displays ingress transport associated primary TP-Tunnels.
<MPLS assoc tp-tunnel>

pw-type <eth-raw | eth-tagged > filters by pseudowire type.

pw-mode <mesh | spoke> filters by pseudowire mode.

service-delimiter-vid <NUMBER: 1-4094> filters by service delimiter VLAN ID.

mtu <#1500..9128> filters by MTU size.

status-tlv <on | off> filters by status TLV settings.

pw-vccv-profile <MPLS Pseudowire VCCV Profile> filters by pseudowire VCCV profile.

pw-cos-profile <MPLS Pseudowire COS Profile> filters by pseudowire COS profile.

To display detailed output of a selected virtual circuit

- 3 Display detailed output of a selected virtual circuit:
 mpls l2-vpn show vc <vc>

To display traffic statistics of a virtual circuit

- 4 Display pseudowires by virtual circuit:
 mpls l2-vpn show vc <vc-name> statistics

To display pseudowires by customer name

- 5 Display pseudowires by customer name:
 mpls l2-vpn show customer <customer>

To display virtual circuit next hops

- 6 Display virtual circuit next hops:
 mpls l2-vpn show vc-nexthops <vc-nexthops>

To display virtual circuit information based on tunnel group

- 7 Display virtual circuit information based on tunnel group:
 mpls l2-vpn show vc-vifs <vc-vifs>

To display virtual circuit resource summary

- 8 Display virtual circuit resource summary:
 mpls l2-vpn show resource summary

To display virtual circuit groups

- 9 Display virtual circuit groups:
 mpls l2-vpn show vc-group
 mpls l2-vpn show vc-group group-id <pw-group-id>
 —end—

Examples

The following example shows the mpls l2-vpn show vc <vc> command:

```
> mpls l2-vpn show vc dvc
```

Pseudowire ID	20001
Customer Name	
Pseudowire Name	dvc
Signaling Type	Dynamic
Pseudowire Admin State	Enabled
Operational State	Up
Pseudowire Blocking	No
Remote Peer IP Address	3.3.3.3
Incoming Label	8197
Outgoing Label	8197
Status TLV	On
Local Fault	None
Remote Fault	None
MTU Size	1500
Pseudowire Type	Ethernet Raw
Pseudowire Mode	Mesh
Forward CoS Profile Name	DefaultPwCoSProfile
Egress L2PT Transform	Disabled
Forward Vccv Profile Name	DefaultPwVccvProfile
Local CC:CV	TTL Gal-GAch : LSP
Remote CC:CV	TTL Gal-GAch : LSP
Operating CC:CV	Gal-GAch : LSP
Virtual Switch Name	data_vs
Failure Reason	n/a
Tunnel Virtual Interface	32769
Tunnel Status	Up
Tunnel Name	uni
Control Word	Off
Operational Control Word	Off

The time for which the pseudowire remains in a forwarding state during switchover is shown in the following example:

```
39XX*> mpls l2-vpn show vc stvc1
```

Pseudowire ID	1
Customer Name	
Pseudowire Name	stvc1

Signaling Type	Static
Pseudowire Admin State	Enabled
Operational State	Up
Pseudowire Blocking	No
Remote Peer IP Address	4.4.4.4
Incoming Label	7001
Outgoing Label	7000
Status TLV	Off
Status Query	On
Refresh Status Interval	600 seconds
Local Fault	n/a
Remote Fault	n/a
PW Bundle Index	1
PW Reversion	On
Reversion hold-down time	30 seconds
PW Protection Role	Primary
PW Protection State	Active
MTU Size	1500
Pseudowire Type	Ethernet Raw
Pseudowire Mode	Mesh
Forward CoS Profile Name	DefaultPwCoSProfile
Egress L2PT Transform	Disabled
Forward Vccv Profile Name	DefaultPwVccvProfile
Local CC:CV	n/a
Remote CC:CV	n/a
Operating CC:CV	Gal-GAch : LSP
Virtual Switch Name	vs1
Failure Reason	n/a
Tunnel Virtual Interface	32769
Tunnel Status	Up
Tunnel Name	st_1
Control Word	Off
Operational Control Word	Off
Forwarding State Up Time	0d 0h 0m 3s

OAM Signalling Statistics	

Status Frames Sent	0
Status Frames Received	1
Status Acks Sent	1
Status Acks Received	0
Status Bad PDUs Received	0
Status Bad Acks Received	0
Mac Withdrawal Frames Sent	1
Mac Withdrawal Frames Received	0
Mac Withdrawal Acks Sent	0
Mac Withdrawal Acks Received	1
Mac Withdrawal Bad PDUs Received	0
Mac Withdrawal Bad Acks Received	0
Mac Withdrawal Sent Sequence Number	1
Mac Withdrawal Received Sequence Number	0

The following example shows traffic statistics for a virtual circuit:

```
5142-140*> mpls 12-vpn show vc ct2TE traffic-statistics
```

-----MPLS VC TRAFFIC STATISTICS-----	
Parameter	Value
VC Name	ct2TE
Rx Packets	17211717
Tx Packets	41375835
Rx Bytes	157074129342

12-20 Virtual circuit configuration

Tx Bytes	377513559790

The following example shows traffic statistics for an S-PE NE node:

```
nim-5160-96*> mpls l2-vpn show vc sCoPw1 statistics
```

-----MPLS VC TRAFFIC STATISTICS-----	
Parameter	Value
VC Name	sCoPw1
Rx Packets	34398
Tx Packets	20750
Rx Bytes	3233412
Tx Bytes	2033500

This example shows the output for a virtual circuit configured with pw-group-id:

```
pns-202-5160*> mpls l2-vpn show vc vc51
```

Pseudowire ID	101

Customer Name	
Pseudowire Name	vc51
Signaling Type	Static
Pseudowire Admin State	Enabled
Operational State	Up
Pseudowire Blocking	No
Remote Peer IP Address	201.201.201.201
Incoming Label	4201
Outgoing Label	5201
Status TLV	On
Group Index	1
Infra Vc	Yes
Status Query	On
Refresh Status Interval	600 seconds
Local Fault	None
Remote Fault	None
MTU Size	1500
Pseudowire Type	Ethernet Raw
Pseudowire Mode	Mesh
Forward CoS Profile Name	DefaultPwCoSProfile
Egress L2PT Transform	Disabled
Forward Vccv Profile Name	vccv-1
Local CC:CV	CW : LSP BFD-ACH-ONLY
Remote CC:CV	n/a
Operating CC:CV	CW : LSP BFD-ACH-ONLY
BFD Monitoring	Enabled
BFD Profile ID	6
BFD Profile Name	vc-bfdProf10
BFD Session ID	33281
BFD Session Name	VBFS_0000000101_vc51
Virtual Switch Name	vs51

Failure Reason	n/a
Tunnel Virtual Interface	32769
Tunnel Status	Up
Tunnel Name	tnl-ingr-202-201-1
Control Word	On
Operational Control Word	On

The following examples of virtual circuit group output list all the pseudowire members and Infra PW belonging to the pseudowire group:

pns-202-5160> mpls l2-vpn show vc-group

VIRTUAL CIRCUIT GROUPS													
Group Index	Primary PW Info			PW Members ID									
	Infra PwId	PW Count											
1	101	105		309, 287, 265, 243, 221, 199, 177, 155, 133, 111,	307, 285, 263, 241, 219, 197, 175, 153, 131, 109,	305, 283, 261, 239, 217, 195, 173, 151, 129, 107,	303, 281, 259, 237, 215, 193, 171, 169, 149, 127,	301, 279, 257, 235, 213, 191, 189, 167, 145, 123,	299, 277, 255, 233, 211, 209, 187, 185, 165, 143, 121,	297, 275, 253, 231, 229, 207, 187, 185, 163, 141, 119,	295, 273, 251, 249, 227, 205, 183, 181, 161, 159, 157,	293, 271, 269, 247, 225, 203, 181, 179, 157, 135,	291, 267, 245, 223, 201, 179, 157, 135, 113,
2	409	51		509, 487, 465, 443, 421,	507, 485, 463, 441, 419,	505, 483, 461, 439, 417,	503, 481, 459, 437, 415,	501, 479, 457, 435, 413,	499, 497, 495, 493, 491,	497, 477, 475, 473, 471,	495, 475, 453, 451, 449,	493, 473, 469, 447, 445,	489, 467, 445, 423, 409

Protection PW Info													
Group Index	PW Info			PW Members ID									
	Infra PwId	PW Count											
1	-	104		310, 288, 266, 244, 222, 200, 178, 156, 134, 112,	308, 286, 264, 242, 220, 198, 176, 154, 132, 110,	306, 284, 262, 240, 218, 196, 174, 152, 130, 108,	304, 282, 260, 238, 216, 194, 172, 150, 128, 106,	302, 280, 258, 236, 214, 192, 170, 168, 146, 124,	300, 278, 256, 234, 212, 190, 188, 166, 144, 122,	298, 276, 254, 232, 210, 208, 186, 164, 142, 120,	296, 274, 252, 230, 208, 206, 184, 162, 140, 118,	294, 272, 250, 228, 206, 204, 182, 160, 138, 116,	292, 270, 248, 226, 202, 180, 158, 136, 114,
2	-	50		510, 488, 466, 444, 422,	508, 486, 464, 442, 420,	506, 484, 462, 440, 418,	504, 482, 460, 438, 416,	502, 480, 458, 436, 414,	500, 478, 456, 434, 412,	498, 476, 454, 432, 410,	496, 474, 452, 430, 408,	494, 472, 450, 428, 406,	492, 470, 448, 426, 404,

pns-202-5160> mpls l2-vpn show vc-group group-id 1

VIRTUAL CIRCUIT GROUPS													
Group Index	Primary PW Info			PW Members ID									
	Infra PwId	PW Count											
1	101	105		309, 287, 265, 243,	307, 285, 263, 241,	305, 283, 261, 239,	303, 281, 259, 237,	301, 279, 257, 235,	299, 277, 255, 233,	297, 275, 253, 231,	295, 273, 251, 229,	293, 271, 269, 227,	291, 267, 245, 223,

12-22 Virtual circuit configuration

			221, 219, 217, 215, 213, 211, 209, 207, 205, 203, 201, 199, 197, 195, 193, 191, 189, 187, 185, 183, 181, 179, 177, 175, 173, 171, 169, 167, 165, 163, 161, 159, 157, 155, 153, 151, 149, 147, 145, 143, 141, 139, 137, 135, 133, 131, 129, 127, 125, 123, 121, 119, 117, 115, 113, 111, 109, 107, 105, 103, 101

Group Index	Infra PwId	PW Count	Protection PW Info PW Members ID
1	-	104	310, 308, 306, 304, 302, 300, 298, 296, 294, 292, 290, 288, 286, 284, 282, 280, 278, 276, 274, 272, 270, 268, 266, 264, 262, 260, 258, 256, 254, 252, 250, 248, 246, 244, 242, 240, 238, 236, 234, 232, 230, 228, 226, 224, 222, 220, 218, 216, 214, 212, 210, 208, 206, 204, 202, 200, 198, 196, 194, 192, 190, 188, 186, 184, 182, 180, 178, 176, 174, 172, 170, 168, 166, 164, 162, 160, 158, 156, 154, 152, 150, 148, 146, 144, 142, 140, 138, 136, 134, 132, 130, 128, 126, 124, 122, 120, 118, 116, 114, 112, 110, 108, 106, 104

Procedure 12-6

Configuring L2-VPN virtual circuit CoS profile

Configure an L2-VPN virtual circuit CoS profile.

Table 12-2 lists the default values for L2-VPN virtual circuit CoS profiles.

Table 12-2
Default values for L2-VPN virtual circuit CoS profiles

Attribute	Default
frame-cos-policy	mapped
frame-cos-map	1 (Map-ID)
fixed-tc	0
resolved-cos-policy	mapped
resolved-cos-map	1 (Map-ID)
resolved-cos-profile	default R-CoS Profile Id
resolved-cos-fixed	0

Step	Action

To create an L2-VPN virtual circuit CoS mapping profile

- 1 Create an L2-VPN virtual circuit CoS mapping profile:

```
mpls l2-vpn pw-cos-profile create cos-profile <cos-profile> [rcos-profile <Resolved CoS Profile Name>]
[resolved-cos-policy <mapped | fixed>]
```

To set configuration parameters for a L2-VPN virtual circuit CoS mapping profile

- 2 Set configuration parameters for a L2-VPN virtual circuit CoS mapping profile:

```
mpls l2-vpn pw-cos-profile set cos-profile <cos-profile>
[rcos-profile <Resolved CoS Profile Name>] [resolved-cos-policy <mapped | fixed>]
```

—end—

Procedure 12-7

Displaying L2-VPN virtual circuit CoS profiles

Display L2-VPN virtual circuit CoS profiles.

Step	Action

To display all L2-VPN virtual circuit CoS profiles

- 1 Display all L2-VPN virtual circuit CoS profiles:

```
mpls l2-vpn pw-cos-profile show
```

To display detailed outputs of a selected L2-VPN virtual circuit CoS profile

- 2 Display detailed outputs of a selected L2-VPN virtual circuit CoS profile:

```
mpls l2-vpn pw-cos-profile show cos-profile <cos-profile-name>
```

—end—

Procedure 12-8

Configuring virtual circuit VCCV profiles

The user can select the CC type 1 (control word) and/or CC type 3 (TTL-exhaust) and/or CC type 4 (GAL/G-ACh or out-of-band-OAM channel). Any combination of these can be selected.

If a virtual circuit connectivity verification (VCCV) profile is not associated with a pseudowire (PW), a default VCCV profile will be selected from the two available default VCCV profiles and this selection will be based on the control word configuration for that PW.

Type 1 or type 3 or type 4 can be used to perform PW segment pings, but the type chosen must be consistent on every segment of the multi-segment PW. As it is desirable to support the PW status and the SP-PE TLV in the MS-PW, CC type 4 is most likely used.

Validity checks are performed at the time of a VCCV profile association with a static PW. The association is rejected if

- No CC-type is enabled
- More than one CC-type is enabled
- Conflicting CC-type is enabled. For example, PW status signaling is enabled, but the VCCV profile has CC-type 3 enabled.
- PW is created having control word configured ON and associate a VCCV profile having CC-3, CC-4.
- PW is created having control word configured OFF and associate a VCCV profile having CC-1, CC-3, CC4.

Step	Action

To create a VCCV profile with the specified parameters**1 Create a VCCV profile:**

```
mpls l2-vpn pw-vccv-profile create vccv-profile <vccv-profile> [cc-in-band <on | off>] {logical-id <NUMBER>} [cc-ttl-exp <on | off>] [cc-ciena-oob <on | off>] [cv-bfd-ach-detect-only <on | off>]
```

where

vccv-profile <vccv-profile>	is the Pseudowire VCCV profile name.
cc-in-band <on off>	enables or disables cc-in-band which supports VCCV CC type 1.
(logical-id <NUMBER>	is the profile index
[cc-ttl-exp <on off>]	sets the TTL-expiry VCCV control channel (CC Type-3) operation.
[cc-ciena-oob <on off>]	sets the Ciena proprietary out-of-band VCCV control channel (CC Type 4) operation.
[cv-bfd-ach-detect-only <on off>]	sets the connectivity verification (CV) type to bit 4 (0x10) (BFD PW-ACH-encapsulated) for pseudowire fault detection only.

To modify a VCCV profile**2 Modify a VCCV profile**

```
mpls l2-vpn pw-vccv-profile set vccv-profile <vccv-profile> [cc-in-band <on|off>] [cc-ttl-exp <on|off>] [cc-ciena-oob <on|off>] [cv-bfd-ach-detect-only <on | off>]
```

where

vccv-profile <vccv-profile>	is the Pseudowire VCCV profile name.
cc-in-band <on off>	enables or disables cc-in-band which supports VCCV CC type 1.
[cc-ttl-exp <on off>]	sets the TTL-expiry VCCV control channel (CC Type-3) operation.
[cc-ciena-oob <on off>]	sets the Ciena proprietary out-of-band VCCV control channel (CC Type 4) operation.
[cv-bfd-ach-detect-only <on off>]	sets the connectivity verification (CV) type to bit 4 (0x10) (BFD PW-ACH-encapsulated) for pseudowire fault detection only.

—end—

Example

The following example shows creation of VCCV profiles:

```
mpls l2-vpn pw-vccv-profile create vccv-profile Profile-CC-3 cc-in-band on cc-  
ciena-oob off cv-bfd-ach-detect-only on
```

```
mpls l2-vpn pw-vccv-profile create vccv-profile Profile-CC-4 cc-in-band on cc-  
ttl-exp off cv-bfd-ach-detect-only on
```

```
mpls l2-vpn pw-vccv-profile create vccv-profile Profile-CC-3-4 cc-in-band on  
cv-bfd-ach-detect-only on
```

```
mpls l2-vpn pw-vccv-profile create vccv-profile Profile-CC-none cc-in-band on  
cc-ttl-exp off cc-ciena-oob off cv-bfd-ach-detect-only on
```

Procedure 12-9

Displaying virtual circuit VCCV profiles

Display the a summary of configured virtual circuit VCCV profile or details about a specific VCCV profile.

Step	Action

1	Display the contents of a virtual circuit VCCV profile: <pre>mpls l2-vpn pw-vccv-profile show [vccv-profile <vccv-profile-name>]</pre>

—end—

Example

The following example shows a list of configured VCCV profiles:

```
> mpls l2 pw-vccv-profile show
```

L2-VPN Pseudowire-VCCV-Profile Table									
VCCV Profile Name	Index	CC IB Type-1	CC OOB Type-2	CC TTL Exp :PW TTL Type-3	CC CIENA OOB Type-4	CV ICMP PING	CV LSP PING	CV BFD PW-ACH Detect Only	USE COUNT
DefaultPwVccvProfile	1	No	No	Yes: 1	Yes	No	Yes	Yes	0
DefaultPwCwVccvProfile	2	Yes	No	Yes: 1	Yes	No	Yes	Yes	0

Procedure 12-10

Configuring MS-PW stitching

Configure MS-PW stitching to create an MS-PW architecture.

Attach virtual circuits to a virtual switch multisegment pseudowire by stitching two PWs together at the S-PE. T-PEs are configured as a single segment. Each PW segment can be statically-configured or dynamically-signaled. However, they can only be statically stitched together at each S-PE.

Step	Action

At the S-PE

- 1 Create the static or dynamic virtual circuit.
For the procedure for configuring a static virtual circuit, refer to [“Configuring static virtual circuits” on page 12-10](#).
For the procedure for configuring a dynamic virtual circuit, refer to [“Configuring dynamic virtual circuits” on page 12-6](#).
- 2 Attach the virtual circuit to the MPLS virtual switch:

```
virtual-switch interface attach mpls-vc <vc_name> vs <vs>
```


where
mpls-vc is the name of the MPLS virtual circuit
<vc_name>
vs <vs> is the name of the MPLS virtual switch

Note: Repeat this step for each virtual circuit you are attaching to a virtual switch.

At the T-PE


- 3 (Optional) Ping the virtual circuit.
For the procedure for pinging the virtual circuit, refer to [“Running ping on a virtual circuit” on page 16-26](#).
- 4 (Optional) Run a traceroute on the virtual circuit:
For the procedure for running a traceroute on the virtual circuit, refer to [“Running a VCCV traceroute” on page 16-17](#).

—end—

Example

The following examples refer to the sample topology shown in [Figure 12-1](#).

Figure 12-1
Multi-segment pseudowire sample topology



Static pseudowire to static pseudowire

The following commands entered at the S-PE create an MS-PW architecture that stitches a static pseudowire named PW1 and a static pseudowire named PW2.

```
mpls l2-vpn create static-vc PW1 pw-id 100 peer 1.1.1.1 in-label 100 out-label 200 pw-cword on tp-tunnel-ingr-corout to-tpe1 pw-type eth-tagged status-tlv on pw-mode switching
mpls l2-vpn create static-vc PW2 pw-id 200 peer 3.3.3.3 in-label 300 out-label 400 pw-cword on tp-tunnel-ingr-corout to-tpe2 pw-type eth-tagged status-tlv on pw-mode switching
virtual-switch create vs S-PE
virtual-switch interface attach mpls-vc PW1 vs S-PE
virtual-switch interface attach mpls-vc PW2 vs S-PE
```

The following command entered at T-PE 1.1.1.1 pings the virtual circuit.

```
mpls ping vc PW1 segment 2 seg-pw-id 333 seg-dest-ip 3.3.3.3 seg-src-ip 2.2.2.2
```

Output for the command is:

```
Legend: '!' - Success, 'X' - Error
! Seq: 1 Latency: 4 ms
! Seq: 2 Latency: 3 ms
! Seq: 3 Latency: 4 ms
! Seq: 4 Latency: 5 ms
! Seq: 5 Latency: 6 ms
```

The following command entered at T-PE 1.1.1.1 runs a traceroute on the virtual circuit.

```
mpls traceroute vc PW segment 2
```

Output for the command is:

```
MPLS traceroute with segment (2) timeout (1000 ms) encap (ip/udp) reply-mode (lsp)
```

Legend: '!' - Success, 'X' - Error
 ! Loc Node: 0 Next S-PE: ip 1.1.1.1 Label 32768 MTU 1500
 ! S-PE: 1 Next S-PE: ip 2.2.2.2 Label 18000, MTU 1500 Latency: 9 ms
 ! Remote T-PE: 2 Latency: 5 ms

Static pseudowire to dynamic pseudowire

The following commands create an MS-PW architecture that stitches a static pseudowire PW1 and a dynamic pseudowire PW2.

```
mpls l2-vpn create static-vc PW1 pw-id 100 peer 1.1.1.1 in-label 100 out-label 200 pw-cword on tp-tunnel-ingr-corout to-tpe1 pw-type eth-tagged status-tlv on pw-mode switching
mpls l2-vpn create dynamic-vc PW2 pw-id 200 peer 3.3.3.3 te-tunnel to-tpe2 pw-cword on pw-type eth-tagged status-tlv on pw-mode switching
virtual-switch create vs S-PE
virtual-switch interface attach mpls-vc PW1 vs S-PE
virtual-switch interface attach mpls-vc PW2 vs S-PE
```

The following command entered at T-PE 1.1.1.1 pings the virtual circuit.

```
mpls ping vc PW1 segment 2 seg-pw-id 333 seg-dest-ip 3.3.3.3 seg-src-ip 2.2.2.2
```

Output for the command is:

Legend: '!' - Success, 'X' - Error
 ! Seq: 1 Latency: 4 ms
 ! Seq: 2 Latency: 3 ms
 ! Seq: 3 Latency: 4 ms
 ! Seq: 4 Latency: 5 ms
 ! Seq: 5 Latency: 6 ms

The following command entered at T-PE 1.1.1.1 runs a traceroute on the virtual circuit.

```
mpls traceroute vc PW segment 2
```

Output for the command is:

```
MPLS traceroute with segment (2) timeout (1000 ms) encap (ip/udp) reply-mode (lsp)
Legend: '!' - Success, 'X' - Error
! Loc Node: 0 Next S-PE: ip 1.1.1.1 Label 32768 MTU 1500
! S-PE: 1 Next S-PE: ip 2.2.2.2 Label 18000, MTU 1500 Latency: 9 ms
! Remote T-PE: 2 Latency: 5 ms
```

Procedure 12-11


Configuring PW-group protection example

Pseudowires can be grouped together for protection by assigning a common “pw-group-id”. The following are PW-group considerations:

- Create Infra PW as the first PW in the group and attach it to a VS before creating any other PWs in the group.
- Infra PW is not allowed with Custom Queue Group.
- Infra PW is not allowed with VC shaping configuration.
- Flush for Infra PW is not allowed through CLI.
- Change of PW-group is not allowed on an Infra PW.
- If a PW member of a PW-group is configured with a custom queue group, it will not participate in the group-level protection and the protection time for that PW would be higher.
- Supported only for Static PW.
- By default, PW are part of group-id 0, which means that they are not grouped. the PW group-id can be set from 1 to 65535.
- The PW group status is controlled by the Infrastructure PW.

The following figure shows the example configuration.

Figure 12-2
PW-group protection configuration example



Step	Action

Configure C1

LSP “tnl-ingr-202-201-1” and “tnl-ingr-202-100-1” already created.
 Type-1 enabled pw-vccv-profile “vccv-1” already created.
 BFD-profile “vc-bfdProf10” already created.

1 Configure static PW and VS for traffic (vs50) and infrastructure (vs51):

```
virtual-switch ethernet create vs vs50
mpls l2-vpn create static-vc vc50 pw-id 99 peer
201.201.201.201 tp-tunnel-ingr-corout tnl-ingr-202-201-1
in-label 4199 out-label 5199 status-tlv on pw-group-id 1
is-infra-pw no
virtual-switch ethernet attach mpls-vc vc50 vs vs50
mpls l2-vpn protection create static-vc vc50-bkp
secondary-pw-id 100 primary-vc-name vc50 peer
100.100.100.100 in-label 4200 out-label 5200 tp-tunnel-
ingr-corout tnl-ingr-202-100-1
virtual-switch ethernet attach mpls-vc vc50-bkp vs vs50
virtual-switch ethernet create vs vs51
mpls l2-vpn create static-vc vc51 pw-id 101 peer
201.201.201.201 tp-tunnel-ingr-corout tnl-ingr-202-201-1
in-label 4201 out-label 5201 status-tlv on pw-vccv-
profile vccv-1 bfd-profile vc-bfdProf10 pw-cword on pw-
group-id 1 is-infra-pw yes
virtual-switch ethernet attach mpls-vc vc51 vs vs51
```

Configure B1

LSP “tnl-ingr-202-201-1” and “tnl-ingr-201-244-1” already created.
 Type-1 enabled pw-vccv-profile “vccv-1” already created.

2 Configure static PW and VS for traffic (vs50) and infrastructure (vs51):

```
virtual-switch ethernet create vs vs50
mpls l2-vpn create static-vc vc50-11 pw-id 99 peer
202.202.202.202 tp-tunnel-egrs-corout-static tnl-egrs-
202-201-1 in-label 5199 out-label 4199 status-tlv on pw-
mode switching is-infra-pw no
virtual-switch ethernet attach mpls-vc vc50-11 vs vs50
mpls l2-vpn create static-vc vc50-12 pw-id 100 peer
24.24.24.24 tp-tunnel-ingr-corout tnl-ingr-201-244-1 in-
label 4199 out-label 5199 status-tlv on pw-mode switching
is-infra-pw no
virtual-switch ethernet attach mpls-vc vc50-12 vs vs50
```

```
virtual-switch ethernet create vs vs51
mpls 12-vpn create static-vc vc51-11 pw-id 101 peer
202.202.202.202 tp-tunnel-egrs-corout-static tnl-egrs-
202-201-1 in-label 5201 out-label 4201 status-tlv on pw-
mode switching pw-vccv-profile vccv-1 pw-cword on is-
infra-pw yes
virtual-switch ethernet attach mpls-vc vc51-11 vs vs51
mpls 12-vpn create static-vc vc51-12 pw-id 102 peer
24.24.24.24 tp-tunnel-ingr-corout tnl-ingr-201-244-1 in-
label 4201 out-label 5201 status-tlv on pw-mode switching
pw-vccv-profile vccv-1 pw-cword on is-infra-pw yes
virtual-switch ethernet attach mpls-vc vc51-12 vs vs51
```

Configure B2

LSP “tnl-egrs-201-244-1” and “tnl-ingr-244-221-1” already created.
Type-1 enabled pw-vccv-profile “vccv-1” already created.

3 Configure static PW and VS for traffic (vs50) and infrastructure (vs51):

```
virtual-switch ethernet create vs vs50
mpls 12-vpn create static-vc vc50-11 pw-id 100 peer
201.201.201.201 tp-tunnel-egrs-corout-static tnl-egrs-
201-244-1 in-label 5199 out-label 4199 status-tlv on pw-
mode switching is-infra-pw no
virtual-switch ethernet attach mpls-vc vc50-11 vs vs50
mpls 12-vpn create static-vc vc50-12 pw-id 99 peer
221.221.221.221 tp-tunnel-ingr-corout tnl-ingr-244-221-1
in-label 4199 out-label 5199 status-tlv on pw-mode
switching is-infra-pw no
virtual-switch ethernet attach mpls-vc vc50-12 vs vs50
virtual-switch ethernet create vs vs51
mpls 12-vpn create static-vc vc51-11 pw-id 102 peer
201.201.201.201 tp-tunnel-egrs-corout-static tnl-egrs-
201-244-1 in-label 5201 out-label 4201 status-tlv on pw-
mode switching pw-vccv-profile vccv-1 pw-cword on is-
infra-pw yes
virtual-switch ethernet attach mpls-vc vc51-11 vs vs51
mpls 12-vpn create static-vc vc51-12 pw-id 101 peer
221.221.221.221 tp-tunnel-ingr-corout tnl-ingr-244-221-1
in-label 4201 out-label 5201 status-tlv on pw-mode
switching pw-vccv-profile vccv-1 pw-cword on is-infra-pw
yes
virtual-switch ethernet attach mpls-vc vc51-12 vs vs51
```

Configure B3

LSP “tnl-ingr-100-67-1” and “tnl-egrs-202-100-1” already created.

4 Configure static PW and VS for traffic (vs50):

```
virtual-switch ethernet create vs vs50
mpls l2-vpn create static-vc vc50-11 pw-id 100 peer
202.202.202.202 tp-tunnel-egrs-corout-static tnl-egrs-
202-100-1 in-label 5200 out-label 4200 status-tlv on pw-
mode switching
virtual-switch ethernet attach mpls-vc vc50-11 vs vs50
mpls l2-vpn create static-vc vc50-12 pw-id 99 peer
67.67.67.67 tp-tunnel-ingr-corout tnl-ingr-100-67-1 in-
label 4200 out-label 5200 status-tlv on pw-mode switching
virtual-switch ethernet attach mpls-vc vc50-12 vs vs50
```

Configure B4

LSP “tnl-egrs-100-67-1” and “tnl-ingr-67-221-1” already created.

5 Configure static PW and VS for traffic (vs50):

```
virtual-switch ethernet create vs vs50
mpls l2-vpn create static-vc vc50-11 pw-id 99 peer
100.100.100.100 tp-tunnel-egrs-corout-static tnl-egrs-
100-67-1 in-label 5200 out-label 4200 status-tlv on pw-
mode switching
virtual-switch ethernet attach mpls-vc vc50-11 vs vs50
mpls l2-vpn create static-vc vc50-12 pw-id 100 peer
221.221.221.221 tp-tunnel-ingr-corout tnl-ingr-67-221-1
in-label 4200 out-label 5200 status-tlv on pw-mode
switching
virtual-switch ethernet attach mpls-vc vc50-12 vs vs50
```

Configure A1

LSP “tnl-egrs-244-221-1” and “tnl-egrs-67-221-1” already created.
 Type-1 enabled pw-vccv-profile “vccv-1” already created.
 BFD-profile “vc-bfdProf10” already created.

6 Configure static PW and VS for traffic (vs50) and infrastructure (vs51):

```
virtual-switch ethernet create vs vs50

mpls l2-vpn create static-vc vc50-11 pw-id 99 peer
24.24.24.24 tp-tunnel-egrs-corout-static tnl-egrs-244-
221-1 in-label 5199 out-label 4199 status-tlv on pw-
group-id 1 is-infra-pw no

virtual-switch ethernet attach mpls-vc vc50-11 vs vs50

mpls l2-vpn create static-vc vc50-12 pw-id 100 peer
67.67.67.67 tp-tunnel-egrs-corout-static tnl-egrs-67-
221-1 in-label 5200 out-label 4200 status-tlv on pw-
group-id 1001 is-infra-pw no

virtual-switch ethernet attach mpls-vc vc50-12 vs vs50

virtual-switch ethernet create vs vs51

mpls l2-vpn create static-vc vc51-11 pw-id 101 peer
24.24.24.24 tp-tunnel-egrs-corout-static tnl-egrs-244-
221-1 in-label 5201 out-label 4201 status-tlv on pw-vccv-
profile vccv-1 bfd-profile vc-bfdProf10 pw-cword on pw-
group-id 1 is-infra-pw yes

virtual-switch ethernet attach mpls-vc vc51-11 vs vs51
```

—end—

Example

The following example shows VC group 1 details:

```
> mpls l2-vpn show vc-group group-id 1
```

VIRTUAL CIRCUIT GROUPS			
Group Index	Infra PwId	Primary PW Info	
		PW Members	ID
1	101	101, 99, 97, 95, 93, 91, 89, 87, 85, 83, 81, 79, 77, 75, 73, 71, 69, 67, 65, 63, 61, 59, 57, 55, 53, 51, 49, 47, 45, 43, 41, 39, 37, 35, 33, 31, 29, 27, 25, 23, 21, 19, 17, 15, 13, 11, 9, 7, 5, 3, 1	
Group Index	Infra PwId	Protection PW Info	
		PW Members	ID
1	-	100, 98, 96, 94, 92, 90, 88, 86, 84, 82, 80, 78, 76, 74, 72, 70, 68, 66, 64, 62, 60, 58, 56, 54, 52, 50, 48, 46, 44, 42, 40, 38, 36, 34, 32, 30, 28, 26, 24, 22, 20, 18, 16, 14, 12, 10, 8, 6, 4, 2	

Virtual switch configuration

Configure features on the virtual switch according to the network plan.

In previous SAOS releases, primary and protected pseudowires (PWs) were allowed to be attached to different virtual switches (VS). This caused traffic switching to be impossible. With SAOS 6.18.1, a check is introduced that does not allow the primary and protected PWs to be attached to different VS. For switching of traffic from primary to protected and vice versa, these PWs must be attached to the same VS.

Procedures

This chapter provides the following procedures:

- [“Configuring Control Frame Tunneling over MPLS” on page 13-3](#)
- [“Attaching virtual circuits to a virtual switch multi-segment pseudowire” on page 13-4](#)
- [“Creating an MPLS management virtual switch” on page 13-5](#)
- [“Changing the management virtual switch” on page 13-7](#)
- [“Configuring multiple VLANs to the MPLS virtual switch” on page 13-8](#)
- [“Creating a VS configuration with UNI only with bundled CVIDs” on page 13-11](#)
- [“Configuring ingress and egress I2-transforms” on page 13-13](#)
- [“Configuring the encapsulation egress I2-transform CoS values” on page 13-17](#)
- [“Configuring untagged frames” on page 13-18](#)
- [“Configuring egress I2-transform CoS values” on page 13-20](#)
- [“Configuring CoS profiles for MPLS pseudowire” on page 13-21](#)
- [“Configuring MEF8 TDM services over MPLS” on page 13-22](#)
- [“Displaying CoS profiles for MPLS tunnels” on page 13-24](#)
- [“Displaying CoS profiles” on page 13-25](#)
- [“Displaying virtual switch information” on page 13-27](#)

13-2 Virtual switch configuration

- [“Displaying MAC table information” on page 13-29](#)
- [“Configuring dual-tag EVPL classification” on page 13-30](#)

Procedure 13-1

Configuring Control Frame Tunneling over MPLS

Configure Control Frame Tunneling (CFT) over MPLS.

Note: Transparent mode is the only mode supported for MPLS.

Step	Action
1	Set layer 2 CFT on the virtual switch to transparent: <pre>virtual-switch l2-cft set vs <vs_name> tunnel-method transparent</pre> where vs <vs_name> is the virtual switch name.
2	Add protocols to layer 2 CFT on the virtual switch: <pre>virtual-switch l2-cft protocol add vs <vs-name> ctrl-protocol lldp disposition forward virtual-switch l2-cft protocol add vs <vs-name> ctrl-protocol lacp disposition discard</pre> where vs <vs_name> is the virtual switch name.
3	Enable layer 2 CFT on the virtual switch: <pre>virtual-switch l2-cft enable vs vpls1</pre> —end—

Example

The following example configures CFT over MPLS.

```
virtual-switch l2-cft set vs vpls1 tunnel-method transparent
virtual-switch l2-cft protocol add vs vpls1 ctrl-protocol lldp disposition forward
virtual-switch l2-cft protocol add vs vpls1 ctrl-protocol lacp disposition discard
virtual-switch l2-cft enable vs vpls1
```

Procedure 13-2

Attaching virtual circuits to a virtual switch multi-segment pseudowire

This procedure consists of stitching two switching pseudowires together on the switching PE. T-PE does not distinguish if a PW is a single PW or part of a multi-segment PW.

Step	Action
1	Create the virtual switch. <pre>virtual-switch ethernet create vs <vs_name> mode vpws</pre> where vs <vs-name> is the name of the MPLS virtual switch
2	Attach the virtual circuit to the MPLS virtual switch: <pre>virtual-switch ethernet attach mpls-vc <vc_name> vs <vs_name> vc <vc_name></pre> where vc <vc-name> is the name of the MPLS virtual circuit vs <vs-name> is the name of the MPLS virtual switch Note: Repeat this step for each virtual circuit you are attaching to a virtual switch. —end—

Example

The following example shows how to attach virtual circuits to the MPLS virtual switch MS-PW.

```
virtual-switch ethernet create vs vs-mspw mode vpws
virtual-switch ethernet attach mpls-vs SVCoS-26-00 vs vs-mspw
virtual-switch ethernet attach mpls-vc SVCoS-21-00 vs vs-mspw
```

Procedure 13-3

Creating an MPLS management virtual switch

Create an MPLS management virtual switch by creating an MPLS virtual switch and associating it with the remote interface.

Note 1: In order to carry remote interface traffic over an MPLS tunnel, the remote-interface is associated with a virtual switch. Management access to the switch can then be gained from any of the members of this virtual switch, including attachment circuit members. Thus, if ACs exist on the virtual switch that is associated with the Remote Interface, customers could obtain management access to the node.

In order to prevent this, create an MPLS virtual switch specifically for use for in-band management, and DO NOT attach customer ACs to that virtual switch.

Note 2: For 3916, 3930 and 3931 platforms ensure that resource allocations have been adjusted. See [“Allocating resources for an MPLS management virtual switch \(3916, 3926m, 3928, 3930 and 3931 platforms\)”](#) on page 15-10.

Note 3: Associating an MPLS virtual switch with the remote interface does not change the basic properties of that virtual switch. Support for features such as Traffic Profiling, L2-CFT, COS Mapping is unchanged.

If the CLI session in use is connected to the remote interface, the user loses access by means of that current session as the remote interface is reconfigured. The user must initiate a new session over the newly-configured virtual switch.

At the node where VLAN-based management transitions to inband MPLS-based management, the VLAN-based management traffic must be added to the MPLS virtual switch as an EVPL attachment circuit, where the VLAN used is the management VLAN.

Note: When an MPLS EPL is added to a UNI port, this UNI port is automatically removed from the default VLAN. For an EVPL, you can remove the UNI Port from the default VLAN by using the CLI. When the default VLAN membership for the EPL/EVPL port has been removed, traffic that is flooded on the default VLAN no longer egresses the UNI port.

Step	Action

13-6 Virtual switch configuration

1 Create the virtual switch.

```
virtual-switch ethernet create vs <vs_name> mode vpls
```

where

vs <vs-name> is the name of the MPLS virtual switch

2 Attach the virtual switch to the MPLS virtual circuit:

```
virtual-switch ethernet attach vs <vs_name> mpls_vc  
<vc_name>
```

where

vs <vs-name> is the name of the MPLS virtual switch

vc <vc-name> is the name of the MPLS virtual circuit

3 Optionally, on 39XX/51XX platforms, add the virtual switch to port and VLAN:

```
virtual-switch ethernet add vs <vs_name> port <n> vlan <n>
```

where

vs <vs-name> is the name of the MPLS virtual switch

4 Associate the virtual switch with the remote interface:

```
interface remote set vs <vs_name>
```

where

vs <vs-name> is the name of the MPLS virtual switch

—end—

Procedure 13-4

Changing the management virtual switch

You can create several MPLS virtual switches, although only one virtual switch can be associated with the remote interface at any point in time.

You can

- replace the management virtual switch with a management VLAN
- replace the management VLAN with a management virtual switch

Step	Action

To replace the management virtual switch with a management VLAN

- 1 Replace the management virtual switch with a management VLAN:

```
interface remote set vlan <VLAN>
```

The old management virtual switch continues to exist and its membership is not changed, but it no longer provides connectivity to the remote interface.

To replace the management VLAN with a management virtual switch

- 2 Replace the management virtual switch with a management VLAN:

```
interface remote set vs <Virtual Switch Name>
```

The old management virtual switch continues to exist and its membership is not changed, but it no longer provides connectivity to the remote interface.

—end—

Procedure 13-5

Configuring multiple VLANs to the MPLS virtual switch

Multiple VLANs can be added or removed to or from a port to the MPLS virtual switch. You can do this by specifying each VLAN or by specifying a range of VLANs.

The following is an example of a VLAN list:

- 100-105
- 100-105, 501-502
- 100
- 100-105, 601

Step	Action

To add or remove multiple VLANs by a specific VLAN

- 1 Add or remove multiple VLANs on the port to the MPLS virtual switch by a specific VLAN:

```
virtual-switch ethernet [add | remove] vs <vs-name>
port <port> vlan <vlan_ID>
```

where

virtual-switch ethernet [add | remove] adds or removes subscriber members to the virtual switch.

vs <vs-name> is the virtual switch name.

port <port> is the port.

vlan <vlan-id> is the VLAN ID.

To add or remove multiple VLANs by a VLAN range

- 2 Add or remove multiple VLANs on the port to the MPLS virtual switch by a VLAN range:

```
virtual-switch ethernet [add | remove] vs <vs-name>
port <port> vlan <vlan-list>
```

where

vs <vs-name> is the virtual switch name.

port <port> is the port.


vlan <vlan-list> is the VLAN range.

—end—

Examples

Figure 13-1 describes CVID and provides an example on how to configure service 1 and service 2. In this example, it configures the services and also attaches a virtual circuit/pseudowire to the virtual switch.

Figure 13-1
CVID bundling



Service 1 S1

```
virtual-switch ethernet create vs svc1 mode vpls
virtual-switch ethernet add vs svc1 port 1 vlan 100
virtual-switch ethernet add vs svc1 port 2 vlan 200
virtual-switch ethernet add vs svc1 port 2 vlan 201
virtual-switch ethernet attach mpls-vc PW1 vs svc1
```

S3

```
virtual-switch ethernet create vs svc1 mode vpls
virtual-switch ethernet add vs svc1 port 1 port S5
virtual-switch ethernet attach mpls-vc PW1 vs svc1
```


Service 2 S1

```
virtual-switch ethernet create vs svc2 mode vpls
virtual-switch ethernet add vs svc2 port 1 vlan 200
virtual-switch ethernet add vs svc2 port 1 vlan 201
virtual-switch ethernet attach mpls-vc PW2 vs svc2
```

S5

```
virtual-switch ethernet create vs svc2 mode vpls
virtual-switch ethernet add vs svc2 port 4
virtual-switch ethernet attach mpls-vc PW2 vs svc2
```

The following example adds multiple VLANs on the port to the MPLS virtual switch.

```
virtual-switch ethernet add vs mpls_vs1 port 1 vlan 101
virtual-switch ethernet add vs mpls_vs1 port 1 vlan 102
virtual-switch ethernet add vs mpls_vs1 port 1 vlan 103
virtual-switch ethernet add vs mpls_vs1 port 1 vlan 105
```

The following example supports one CVID for each line as shown.

```
virtual-switch ethernet add vs vs-mpls_tag port 24 vlan 100
virtual-switch ethernet add vs vs-mpls_tag port 24 vlan 101
virtual-switch ethernet add vs vs-mpls_tag port 24 vlan 103
virtual-switch ethernet add vs vs-mpls_tag port 24 vlan 105
```

Procedure 13-6

Creating a VS configuration with UNI only with bundled CVIDs

Create a VS configuration with UNI only with bundled CVIDs.

Step	Action
1	Set the virtual switch layer 2 transform action to support VLAN translation for virtual switches on the first port: <pre>port set port <PortNameList> vs-l2-transform i-push,e-pop</pre> where <PortNameList> is the name of the first port
2	Set the virtual switch layer 2 transform action to support VLAN translation for virtual switches on the second port: <pre>port set port <PortNameList> vs-l2-transform i-push,e-pop</pre> where <PortNameList> is the name of the second port
3	Reserve VLANs for use in creating Virtual Switch Instances: <pre>virtual-switch add reserved-vlan <VLAN IDList></pre>
4	Create a virtual switch instance: <pre>virtual-switch ethernet create vs <VirtualSwitchName[15]> reserved-vlan <VLAN IDList></pre>
5	Disable the LLDP on the UNI port. <pre>lldp set port <port> disable</pre> where port <port> is the port number.
6	Add attachment circuits for the first port: <pre>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>]</pre>
7	Add attachment circuits for the second port: <pre>virtual-switch ethernet add vs <vs> {port <port>} {vlan <VLAN list>} {ip-interface <interface-name>} {encap-cos-policy <dot1dpri-inherit fixed ip-prec-inherit </pre>

13-12 Virtual switch configuration


```
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

The example illustrated in [Figure 13-2](#) shows a VS configuration where traffic is classified to CVID 555 and 600. The SVID 4000 is pushed on the frame at ingress and popped at egress.

Figure 13-2

VS L2 transform: i-push, e-pop for UNI only with bundled CVIDs



The following example creates the configuration.

```
> port set port 1 vs-l2-transform i-push,e-pop  
> port set port 2 vs-l2-transform i-push,e-pop  
> virtual-switch add reserved-vlan 4000  
> virtual-switch ethernet create vs vs100 reserved-vlan 4000  
> lldp set port 1 disable  
> lldp set port 2 disable  
> virtual-switch ethernet add vs vs100 port 1 vlan 555  
> virtual-switch ethernet add vs vs100 port 1 vlan 600  
> virtual-switch ethernet add vs vs100 port 2 vlan 555  
> virtual-switch ethernet add vs vs100 port 2 vlan 600
```

Procedure 13-7 Configuring ingress and egress L2-transforms

Figure 13-3 shows the ingress-L2-transform and egress-L2-transforms on MPLS EVPL ACs.

Figure 13-3
L2-transforms on ACs




Table 13-1 lists I2-transform capabilities.

Table 13-1
Ingress and egress transform values

Feature	5142, 5160	3926m, 3928, 3942	39f16, 3930, 3931, 3932	5150
Egress-I2-transform				
Pop only for traffic with TPID=8100	supported	supported	not supported	not supported
Push-8100/88a8/9100.<vlanid> for traffic with TPID=8100, 88a8, 9100	supported	supported	supported	supported
Stamp-8100.<vlanid> only, for traffic with TPID=8100	supported	supported	not supported	not supported
Ingress-I2-transform				
Pop for traffic with TPID=8100, 88a8, 9100	supported	supported	supported	supported
Push-8100.<vlanid> only, for traffic with TPID=8100, 88a8, 9100	supported	supported	supported	supported
Stamp-8100.<vlanid> only, for traffic with TPID=8100, 88a8, 9100	supported	supported	supported	supported

For CVID bundling, ingress-I2-transform and egress-I2-transform attributes are required on the bundle and not necessary on per CVIC. This means that the same I2-transform is applied for all VLANs in the bundle, including the untagged data.

Step	Action

To configure the ingress I-2 transform**1** Configure the ingress-I2-transform:

```
virtual-switch ethernet [add | set] vs <vs-name> port
<port> [vlan <vlan>] set ingress-l2-transform <format-
string>
```

where

virtual-switch ethernet [add | set] adds the subscriber members to the virtual switch or sets the Ethernet virtual switch.

vs <vs-name> is the name of the virtual switch.

port <port> is the port on the ingress-I2-transform

vlan <vlan> is the subscriber VLAN ID.

set-ingress-I2-transform <format-string> sets the ingress-I2-transform to either no-op, push or pop. Default is no-op.

To configure the egress I-2 transform**2** Configure the egress-I2-transform:

```
virtual-switch ethernet [add | set] vs <vs-name> port
<port> [vlan <vlan>] set egress-l2-transform <format-
string>
```

where

virtual-switch ethernet [add | set] adds the subscriber members to the virtual switch or sets the Ethernet virtual switch.

vs <vs-name> is the name of the virtual switch.

port <port> is the port on the egress-I2-transform

vlan <vlan> is the VLAN.

set-egress-I2-transform <format-string> sets the egress-I2-transform. Default is no-op.

—end—

Example

The following command configures the ingress-I2-transform on MPLS EVPL ACs:

```
virtual-switch ethernet add vs10 port 10 vlan vlan15 ingress-l2-transform
stamp-*.100
```

The following command configures the egress-I2-transform on MPLS EVPL ACs:

13-16 Virtual switch configuration

```
virtual-switch ethernet add vs10 port 10 vlan vlan15 egress-l2-transform pop
```

Procedure 13-8

Configuring the encapsulation egress I2-transform CoS values

Configure the encapsulation ingress I2-transform CoS fixed/mapped values.

Note: This procedure is supported on the 3916, 3926m, 3928, 3930, 3931, 3932, 3942, 5142 and 5160 platforms.

Step	Action

1	Configure the ingress-I2-transform CoS values:

```
virtual-switch ethernet set port <port-id> vs <vs>
vlan <vlan-id> encap-cos-policy <dot1dpri-inherit |
fixed> encap-fixed-dot1dpri <NUMBER: 0..7>
```

where

port <port-id> is the subscriber port.

vs <vs> is the virtual switch Ethernet name.

vlan <vlan-id> is the VLAN

encap-cos-policy <dot1dpri-inherit | fixed> sets the encap-cos-policy to fixed dot1dpri.

encap-fixed-dot1dpri <NUMBER: 0..7> sets the encapsulation fixed-dot1dpri value.

—end—

Example

The following example configures the encapsulation egress I2-transform CoS values:

```
virtual-switch ethernet set port 10 vs mcLag3930-5150 vlan vlan10 encap-cos-
policy dot1dpri-inherit encap-fixed-dot1dpri 7
```


Procedure 13-9

Configuring untagged frames

Configure untagged data or untagged control frames.

Step	Action

To configure an untagged data frame on a virtual switch

- 1 Configure an untagged data frame:


```
port <set | unset> port <PortName> untagged-data-vs
<VirtualSwitchName>
```

where

port <set unset>	sets or unsets the port on the untagged data frame.
port <PortName>	is the name of the port.
untagged-data-vs <VirtualSwitchName>	is the MPLS virtual switch name.

To configure acceptable frame types

- 2 Configure acceptable frame types:


```
port <set | unset> port <PortName> acceptable-frame-
type <all | tagged-only | untagged-only>
```

where

port <set unset>	sets or unsets the port attributes.
port <PortName>	is the name of the port.
acceptable-frame-type <all tagged-only untagged-only>	sets the acceptable-frame-type as all or tagged only.

To configure an untagged control frame on a virtual switch

- 3 Configure an untagged control frame:


```
port <set | unset> port <PortName> untagged-ctrl-vs
<VirtualSwitchName>
```

where

port <set unset>	sets or unsets the port on the untagged control frame.
port <PortName>	is the name of the port.
untagged-ctrl-vs <VirtualSwitchName>	is the MPLS virtual switch name.

—end—

Example

The following example configures an untagged data on a virtual switch.

```
port set port 10 untagged-data-vs mcLag3930-5150
```

The following example configures CFT over MPLS.

```
virtual-switch l2-cft set vs vpls1 tunnel-method transparent
virtual-switch l2-cft protocol add vs vpls1 ctrl-protocol lldp disposition
forward
virtual-switch l2-cft protocol add vs vpls1 ctrl-protocol lacp disposition
discard
virtual-switch l2-cft enable vs vpls1
```

Procedure 13-10

Configuring egress I2-transform CoS values

The resolved CoS (r-CoS) to Frame CoS (F-CoS) map (frame-cos-map) with frame CoS mapping, the R-CoS and R-Color associated with a frame is mapped to specific CoS values in the frame based upon the default frame CoS mapping table (DefaultRcosFcos).

You can also configure customized maps and assign them to ports. You can configure up to three custom frame CoS maps to use the frame's R-CoS and R-Color to map the customer PCP/L2 CoS 802.1D values upon egress.

Note: This procedure is supported on the 3916, 3926m, 3928, 3930, 3931, 3932, 3942, 5142 and 5160.

Step	Action
1	Configure the egress-I2-transform CoS values: <pre>port set port <PortName> frame-cos-map <Frame CoS Map Name></pre> where port <PortName> are the ports to set. frame-cos-map is the name of the Frame CoS map. <Frame CoS Map Name>
2	Set the traffic-services cos-mapping: <pre>traffic-services cos-mapping frame-cos-map set cos-map <cos-map> [dot1dpri-cos <NUMBER: 0..7>]</pre> where cos-map <cos-map> sets the CoS map. dot1dpri-cos sets the dot1dpri-cos value. <NUMBER: 0..7> —end—

Example

The following example configures egress I-2 transform CoS values:

```
port set port 10 frame-cos-map DefaultRcosFcos
traffic-services cos-mapping frame-cos-map set cos-map DefaultRcosFcos
dot1dpri-cos 7
```

Procedure 13-11

Configuring CoS profiles for MPLS pseudowire

Configure CoS profiles for MPLS pseudowire (PW).

Note: This procedure is supported on the 3916, 3926m, 3928, 3930, 3931, 3932, 3942, 5142 and 5160.

Step	Action
1	Configure CoS profiles for MPLS PW <pre>mpls l2-vpn create [static-vc dynamic-vc] [pw-cosprofile <MPLS Pseudowire COS Profile></pre> where pw-cosprofile <MPLS is the PW-cos profile. Pseudowire COS Profile>

—end—

Example

The following example configures CoS profiles for MPLS pseudowire:

```
mpls l2-vpn create dynamic-vc 15 pw-cos-profile DefaultPwCosProfile
```

Procedure 13-12

Configuring MEF8 TDM services over MPLS

Configure MEF8 TDM services over MPLS as required by the network plan.

Step	Action
1	Set the TDM port mode to ANSI: <code>port tdm set mode ansi</code>
2	Create a virtual switch with a mode of VPWS: <code>virtual-switch ethernet create vs <vs> mode vpws</code>
3	Create a TDM profile: <code>virtual-circuit tdm create tdm-profile <tdm-profile> pwType satop payload-size <BYTES></code>
4	Create a static MPLS virtual circuit: <code>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> tp-tunnel-egrs- corout-dynamic <tunnel name></code>
5	Attach members to an MPLS capable virtual switch: <code>virtual-switch ethernet attach mpls-vc <mpls-vc> vs <VS></code>
6	Enable the TDM port: <code>port tdm enable port <port></code>
7	Create a satop TDM attachment circuit: <code>attachment-circuit satop create ac <ac> port <port></code>
8	Create TDM virtual circuit: <code>virtual-circuit tdm-mef8 create vc <vc> ac <ac> tdm- profile <TDM profile name> peer-mac <MAC address: XX:XX:XX:XX:XX:XX> c-vid <NUMBER: 1-4094> c-pcp <NUMBER: 0-7> in-ecn <NUMBER: 1-65535> out-ecn <NUMBER: 1-65535></code>
9	Add subscriber members to the virtual switch: <code>virtual-switch ethernet add vs <vs> tdm-vc <tdm-vc> vlan <vlan></code>

—end—

Example

The following example configures a MEF8 TDM service over MPLS.

```
port tdm set mode ansi
virtual-switch ethernet create vs TDMtest mode vpws
virtual-circuit tdm create tdm-profile satop pwType satop payload-size 192
```

```
mpls l2-vpn create static-vc TDMtest pw-id 100 peer 10.10.10.10 in-label 8000
out-label 8000 tp-tunnel-egrs-corout-dynamic tnl-55-54
virtual-switch ethernet attach mpls-vc TDMtest vs TDMtest
port tdm enable port tdm01
attachment-circuit satop create ac DS1p1 port tdm01
virtual-circuit tdm-mef8 create vc DS1p1 ac DS1p1 tdm-profile satop peer-mac
00:23:8a:ad:59:9d c-vid 222 c-pcp 6 in-ecn 100 out-ecn 100
virtual-switch ethernet add vs TDMtest tdm-vc DS1p1 vlan 100
```

Procedure 13-13

Displaying CoS profiles for MPLS tunnels

Display a summary of CoS profiles for MPLS tunnels.

Step	Action

1	Display a summary of CoS profiles. (option <pre>mpls tunnel-cos-profile show cos-profile <cos-profile></pre> —end—

Example

The following example shows sample output for the mpls tunnel-cos-profile show command.

```
mpls tunnel-cos-profile show
```

MPLS Tunnel-COS-Profile Table								
CoS-Mapping Profile Name	Index	FCoSPolicy	FCoSMapID	FixedTc	RCoSPolicy	RCoSMapID	RCoSFixed	UseCnt
DefaultTunlCoSProfile	1	mapped	1	0	mapped	1	0	8
TE_TUN_COS_PROF	2	fixed	1	4	fixed	1	6	0

Procedure 13-14

Displaying CoS profiles

You can display the

- resolved CoS map
- frame CoS map
- MPLS tunnel cos-profile
- MPLS l2-vpn pw cos-profile

Note: This procedure is supported on the 3916, 3926m, 3928, 3930, 3931, 3932, 3942, 5142 and 5160.

Step	Action

To display the resolved CoS map

- 1 Display the resolved Cos map:

```
traffic-services cos-mapping resolved-cos-map show
cos-map DefaultFcosRcos
```

To display the frame CoS map

- 2 Display the frame-cos-map:

```
traffic-services cos-mapping frame-cos-map show cos-
map DefaultRcosFcos
```

To display the MPLS tunnel cos-profile

- 3 Display the MPLS tunnel cos-profile:

```
mpls tunnel-cos-profile show cos-profile
DefaultTun1CoSProfile
```

To display the MPLS l2-vpn pw cos-profile

- 4 Display the MPLS l2-vpn pw-cos-profile:

```
mpls l2-vpn pw-cos-profile show cos-profile
DefaultPwCosProfile
```

—end—

Example

The following example shows the output for the traffic-services cos-mapping frame-cos-map show cos-map DefaultRcosFcos command:

13-26 Virtual switch configuration

RESOLVED COS MAP TO FRAME COS INFO					
Name logical-id		DefaultRcosFcos 1			
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

Procedure 13-15

Displaying virtual switch information

You can display virtual switch information.

Note: This procedure is supported on the 3916, 3926m, 3928, 3930, 3931, 3932, 3942, 5142 and 5160 platforms.

Step	Action

1	Display virtual switch information: <pre>virtual-switch ethernet show vs <vs-name></pre> where vs <vs_name> is the virtual switch name.

—end—

Example

The following example shows the output for the virtual-switch ethernet show vs sPWsCoTp1-15-18 command:

```
virtual-switch ethernet show vs sPWsCoTp1-15-18
```

MPLS VIRTUAL SWITCH INFO				
Parameter		Value		
Name		sPWsCoTp1-15-18		
ID		1		
Description				
Mode		vpls		
Creation		Static		
CPT Status		Enabled		
CPT Method		transparent		
Transparent Validate		off		
Mac Learning		Enabled		
Subscr Dot1D Policy		leave		
Decap TTL Policy		NA		
Decap CoS Policy		leave		
Decap Fixed Dot1DPri		NA		
Encap CoS Policy		fixed		
Encap Fixed Dot1DPri		2		
Attachment Circuits				
Port	Vlan	Encap Cos Policy	Encap Fixed Dot1DPri	Tpid
1	10	fixed	4	8100
1	11	vs-inherit	2	8100
1	12	vs-inherit	2	8100
1	13	vs-inherit	2	8100
1	14	vs-inherit	2	8100
1	15	vs-inherit	2	8100
Attachment Circuits				
Port	Vlan	Ingress L2 Transform	Egress L2 Transform	

13-28 Virtual switch configuration

1	10	pop	push-9100.333
1	11	pop	push-9100.333
1	12	pop	push-9100.333
1	13	pop	push-9100.333
1	14	pop	push-9100.333
1	15	pop	push-9100.333
----- Virtual Circuits -----			
Name	Decap Label	Encap Label	
sPWsCoTp1-15-18	6001	6001	

Procedure 13-16

Displaying MAC table information

You can display MAC table information.

Step	Action

1	Display MAC table information:

```
flow mac-addr show
```

```
—end—
```

Example

The following example shows the output for the flow mac-addr show command:

```
> flow mac-addr show
Getting MAC entries, please be patient...
```

FLOW MAC-LEARN TABLE				
VLAN	Address (SA)	Port	Type	Virtual Switch
32	9C:7A:03:7E:39:E1	2	Dyna	---
127	9C:7A:03:7E:39:FF	2	Dyna	---
---	00:00:00:00:00:AA	5	Dyna	vs1
---	00:00:00:00:00:55	5	Dyna	vs1
---	00:00:00:00:00:CC	vc1	Dyna	vs1

Procedure 13-17

Configuring dual-tag EVPL classification

Dual-tag classification supports dual-tag EVPL bundling on AC of MPLS VS. The configuration model for dual-tag classification is straightforward and plugged into the existing single-tagged EVPL configuration model.

Prior to SAOS 6.18, only the port and outer VLAN combination was allowed to be added to MPLS VS. With SAOS 6.18, the port, outer VLAN, and inner VLAN combination can be added as AC to MPLS VS. This dual-tag classification is achieved with a new attribute, “vtag-stack”.

Note: To achieve dual-tag classification, the inner TPID of the incoming dual-tag frame should be 0x8100 (by default). The inner TPID can be modified using this CLI command:

```
> vlan set inner-tpid <88A8 | 8100 | 9100>
> vlan enable inner-tpid
```

RCOS/RCOLOR behavior for dual-tag EVPL

In dual-tag classification, for encap-cos-policy port-inherited, RCOS and RCOLOR are derived from the inner VLAN tag of the frame.

For example:

```
> virtual-switch ethernet add vs Mpls_VS port 1 vtag-stack
100:200 encap-cos-policy port-inherit
```

Limitations

Limitations of dual-tag EVPL classification include:

- Hybrid CVID bundling (combination of single- and dual-tag classification) is not supported in the same MPLS VS. All EVPL members must be double-tagged. In the following example, dual-tag and single-tag EVPL are added as part of CVID bundling, which is not supported:

```
> virtual-switch ethernet add vs Mpls_VS port 1 vtag-stack
100:200
> virtual-switch ethernet add vs Mpls_VS port 1 vlan 200
```

- DHCP L2 over MPLS is not supported for dual-tag EVPL bundling.
- CFMoMPLS uses the inner VLAN for “up” MEP creation and the outer VLAN is ignored.
- Traffic profiling (standard TP with port and VLAN-based) is not supported for dual-tag EVPL’s. However, VS-based TP works as expected.
- Benchmark functionality is not supported in dual-tag EVPL.

- IGMP over MPLS is not supported for dual-tag EVPL bundling.

Step	Action

1	Configure dual-tag EVPL classification: <pre>> virtual-switch ethernet add vs <vs-name> port <port-no> vtag-stack <OuterVlan:InnerVlan></pre> where vs <vs-name> is the name of the virtual switch. port <port-no> is the number of the port. vtag-stack is the number of the outer VLAN followed by the <OuterVlan: number of the inner VLAN in the VTag stack. InnerVlan>

—end—

Example

Here is some sample dual-tag EVPL configuration:

```
> virtual-switch ethernet add vs Mpls_VS port 1 vtag-stack 100:200
> virtual-switch ethernet add vs Mpls_VS port 1 vtag-stack 150:200
```

L2 VPN service configuration

The 39XX/51XX Service Delivery, Aggregation and Virtualization Switches support the simultaneous deployment of MPLS-based L2 VPN services, including:

- Virtual Private Wire Services (VPWS)
- Virtual Private Local Area Network Services (VPLS)
- [“Hierarchical VPLS” on page 2-7](#) (H-VPLS)

This chapter provides the following procedures:

- [“Configuring a 39XX/51XX LSR” on page 14-2](#)
- [“Configuring a 39XX/51XX VPWS” on page 14-5](#)
- [“Configuring a 39XX/51XX VPLS” on page 14-9](#)
- [“Configuring a 39XX/51XX H-VPLS” on page 14-13](#)

Procedure 14-1

Configuring a 39XX/51XX LSR


Configure a 39XX/51XX LSR to establish a dynamic base MPLS configuration on a 39XX/51XX switch. The base configuration allows the switch to function as an LSR and builds the foundation for LER functionality. This procedure sets up the L3 interfaces, IGP, and signaling protocols used by MPLS.

For SAOS 6.13 and higher, the IP interface Maximum Transmission Unit (MTU) has to be set appropriately in order for the MPLS data traffic with a packet size greater than 1500 bytes to pass through. The default MTU value is 1500 which means any data traffic greater than 1500 bytes will be dropped. The maximum MTU value that can be set is 9216.

Note: This IP interface MTU setting is applicable for devices running SAOS 6.13 and higher.

Figure 14-1 shows the sample physical topology.

Figure 14-1
Sample physical topology



Step	Action

1	Create a VLAN for each L3 interface / physical port pair. <pre>vlan create vlan <vlan_#></pre>

- 2 Add the L3 Interface VLAN to the associated physical port.

```
vlan add vlan <vlan_#> port <port_#>
```
- 3 If the switch is part of a ring topology, remove VLAN 1,127 from each of the physical ports.

```
vlan remove vlan 1,127 port <port_#>
```
- 4 Disable RSTP.

```
rstp disable
```
- 5 Create the L3 and loopback interfaces.

```
interface create loopback <name> ip <ip>
interface create ip-interface <name> ip <ip> subnet
<subnet> [vlan <vlan>|vs <vs-name>] ip-forwarding
<on|off> if-num <NUMBER: 1-4294967295> [priority
<NUMBER:0-7>]
```
- 6 Set the MTU size if your device is running SAOS 6.13:

```
interface set ip-interface <name> mtu <NUMBER:
1500..9216>
```
- 7 Create an IGP instance.

```
isis instance create isis-instance <instance_name> level
L1 area <area_#>
```
- 8 Attach each L3 interface to the IGP instance.

```
isis interface attach ip-interface <name> isis-instance
<name> level L1
```
- 9 Enable RSVP globally and on each of the interfaces previously created.

```
rsvp-te enable
rsvp-te enable ip-interface <name>
```
- 10 Enable LDP globally.

```
ldp enable
```

—end—

Example

The following sample configuration file segment shows the base LSR configuration for the 3930 named R3 in the network topology diagram shown in [Figure 14-1 on page 14-2](#).

```
!-----
!      General Device Config
!-----
system set host-name R3

! VLANs 1,127 have been removed from all ports in the base configuration. This
! eliminates the possibility of L2 loops which allows RSTP to be disabled
! globally.
rstp disable
interface create loopback lb ip 3.3.3.3
isis instance create isis-instance region1 level L1 area 46.0001
```

14-4 L2 VPN service configuration

```
!-----
!Interface if104
!-----
!A different VLAN is used for each L3 Interface / physical port combination to prevent a flood domain from forming.
vlan create vlan 104
vlan add vlan 104 port 9
vlan remove vlan 1,127 port 9

!The L3 interface named if104 is associated with VLAN 104 to coincide with port 9 above. IP Forwarding is also enabled to allow routing between L3 interfaces. The interface identifier for a non-IP MPLS link is also specified.
interface create ip-interface if104 ip 10.104.13.3 subnet 255.255.255.0 vlan 104 ip-forwarding on if-num 25000

!Each L3 interface is attached to the IGP instance to signal that it should be included in routing updates and to enable IGP signaling for that link.
isis interface attach ip-interface if104 isis-instance region1 level L1

!RSVP provides the signaling and label exchange mechanism that will be used to generate the LSP. It must be enabled both globally and on each interface.
rsvp-te enable ip-interface loopback
rsvp-te enable
rsvp-te enable ip-interface if104

!LDP provides the signaling and label exchange mechanism that will be used to generate the Pseudowire. LDP is enabled globally.
ldp enable

!-----
!Interface if103
!-----
vlan create vlan 103
vlan add vlan 103 port 10
vlan remove vlan 1,127 port 10

interface create ip-interface if103 ip 10.103.34.3 subnet 255.255.255.0 vlan 103 ip-forwarding on if-num 25000

isis interface attach ip-interface if103 isis-instance region1 level L1

rsvp-te enable ip-interface if103
```

Procedure 14-2

Configuring a 39XX/51XX VPWS


VPWS provides point-to-point connectivity between two remote Local Area Networks (LANs). In this example protection switching is performed through tunnel redundancy.

FEC-129 is supported for mesh PW only.

Note: This procedure assumes that the device has already been configured with base LSR functionality as illustrated in “Configuring a 39XX/51XX LSR” on page 14-2.

Figure 14-2 shows the VPWS topology.

Figure 14-2
VPWS Topology



Step	Action
1	Create a virtual switch. <pre>virtual-switch ethernet create vs <vs> mode <vpws vpls></pre> where vs <vs> is the virtual switch name. mode <vpwslvpls> is the capability of the virtual-switch.
2	Disable the LLDP on the UNI port. <pre>lldp set port <port> disable</pre> where port <port> is the port number.
3	Attach the UNI port(s) to the virtual switch. <pre>virtual-switch ethernet add vs <vs> port <Port list> [vlan <VLAN list>]</pre> where vs <vs> is the virtual switch name. port <Port list> is the subscriber port list. vlan <VLAN list> is the subscriber VLAN ID.
4	Create a primary and backup path. <pre>rsvp-te path create rsvp-path <rsvp-path></pre> where rsvp-path <rsvp-path> is the RSVP path to create. <pre>rsvp-te path set rsvp-path <rsvp-path> index <NUMBER: 1-16> ip <IP address> hop-type <strict loose></pre> where rsvp-path <rsvp-path> is the RSVP path to create. index <NUMBER: 1-16> is the position of the node in the rsvp path list. ip <IP address> is the IP address. hop-type <strict loose> is the hop type.

5 Create a primary and backup tunnel and apply the previously created path.

```
mpls tunnel create rsvp-ingress <rsvp-ingress> dest-ip
<IP address> explicit-tunnel-path <MPLS Rsvp Path>
[backup-tunnel <MPLS ingress primary tunnel>]
```

where

rsvp-ingress is the RSVP tunnel name.
 <rsvp-ingress>
 dest-ip <IP address> is the tunnel destination IP address.
 path <MPLS Rsvp Path> is the RSVP path name.
 [backup-tunnel <MPLS ingress primary tunnel>] is the list of primary tunnels.

6 Create a virtual circuit.

```
mpls l2-vpn create dynamic-vc <dynamic-vc> pw-id <NUMBER:
1-2147483647> peer <IP address> tunnel <MPLS ingress
primary tunnel> pw-type <eth-raw|eth-tagged> pw-mode
<mesh|spoke> fec-129
```

where

dynamic-vc is the dynamic virtual switch name.
 <dynamic-vc>
 pw-id <NUMBER: 1-2147483647> is the VPN ID.
 peer <IP address> is the destination IP address.
 tunnel <MPLS ingress primary tunnel> is the transport tunnel.
 pw-type <eth-raw|eth-tagged> is the pseudowire type.
 pw-mode <mesh|spoke> is the pseudowire mode.
 fec-<128|129> configures the pseudowire for mesh PW only. FEC-129 is only supported for mesh PW.

7 Attach the virtual circuit to the virtual switch.

```
virtual-switch ethernet attach mpls-vc <mpls-vc> vs <vs>
```

where

mpls-vc is the list of MPLS VCs.
 <mpls-vc>
 vs <vs> is the virtual switch name.

Example

The following sample configuration file segment shows the VPWS configuration for the 3930 named R1 in the network topology diagram shown in [Figure 14-2 on page 14-5](#).

```

!----- vpws-r1.r4-----!
! ---The primary path is created towards R2. This is also known as an
Explicit Route Object and is used to override the behavior of the IGP.
rsvp-te path create rsvp-path vpws-r1.r4.pri
rsvp-te path set rsvp-path vpws-r1.r4.pri index 5 ip 10.101.12.2 hop-type
loose

! ---Backup Path
rsvp-te path create rsvp-path vpws-r1.r4.bak
rsvp-te path set rsvp-path vpws-r1.r4.bak index 5 ip 10.104.13.3 hop-type
loose

! ---The primary tunnel is created as a dynamic RSVP tunnel. The
destination IP address is targeted at the loopback interface to ensure it
is always up.
mpls tunnel create rsvp-ingress tun-vpws-r1.r4.pri dest-ip 4.4.4.4 explicit-
tunnel-path vpws-r1.r4.pri record-route on
! ---The backup tunnel is associated with the primary tunnel and is only
forwarding if the primary tunnel fails.
mpls tunnel create rsvp-ingress tun-vpws-r1.r4.bak dest-ip 4.4.4.4 explicit-
tunnel-path vpws-r1.r4.bak backup tun-vpws-r1.r4.pri

! --- Virtual Circuit, Virtual Switch, and Attachment Circuit
mpls l2-vpn create dynamic-vc vpws-r1.r4 pw-id 100 peer 4.4.4.4 tunnel tun-
vpws-r1.r4.pri pw-mode mesh pw-type eth-raw fec-129

! --- The Virtual Switch is configured in VPWS mode to prevent more
than 1 attachment circuit from being added. This adheres to the strict
definition of VPWS only having a single attachment circuit.
virtual-switch ethernet create vs vpws-r1.r4 mode vpws
lldp set port 10 disable
virtual-switch ethernet add vs vpws-r1.r4 port 10 vlan 100
virtual-switch ethernet attach mpls-vc vpws-r1.r4 vs vpws-r1.r4

```

Procedure 14-3

Configuring a 39XX/51XX VPLS

VPLS provides point-to-multipoint inter-LAN connectivity. With VPLS, PEs are connected to each other with a full mesh of virtual circuits for each VPLS instance. Each PE connects to an MTU-s on the UNI side through an MPLS network, or it connects to a CE device through an untagged interface or 802.1Q interface.

Note: In Release 6.9, the system handles the outer VLAN which could be C-VLAN, S-VLAN or B-VLAN and not the combination of multiple VLANs as in Q-in-Q or MAC-in-MAC. This is true for a UNI port on a PE also.


Each mesh virtual circuit connects to a remote PE that is a member of the same VPLS instance which provides emulated LAN service for the broadcast domain of the virtual switch. When configuring the mesh virtual circuit, include the virtual circuit identifier. For VPLS configuration, all the mesh virtual circuits that are associated with the same virtual switch may or may not have the same value.

Note 1: In order to be interoperable with other vendors, configure the same virtual circuit identifier for the mesh virtual circuits that are added to the same virtual circuit.

Note 2: This procedure assumes that the device has already been configured with base LSR functionality as illustrated in [“Configuring a 39XX/51XX LSR” on page 14-2](#).

Figure 14-3 on page 14-10 shows the VPLS topology.

Figure 14-3
VPLS topology



Step	Action
1	Create a virtual switch. <pre>virtual-switch ethernet create vs <vs-name> mode <vpws vpls></pre>
2	Add the UNI port(s) to the virtual switch. <pre>virtual-switch ethernet add vs <vs-name> port <name> [vlan <vlan-id>]</pre>
3	Create an explicit path for each PE in the VPLS mesh. <pre>rsvp-te path create rsvp-path <name> rsvp-te path set rsvp-path <name> index <#> ip <ip> hop- type <loose strict></pre>
4	Create an MPLS tunnel for each PE in the MPLS mesh. <pre>mpls tunnel create rsvp-ingress <name> dest-ip <ip> explicit-tunnel-path <name> record-route <on off></pre>
5	Create a virtual circuit (Pseudowire) for each PE in the MPLS mesh. <pre>mpls l2-vpn create dynamic-vc <name> pw-id <#> peer <ip> tunnel <name> pw-mode <spoke mesh> pw-type <eth-raw eth- tagged></pre>
6	Attach each virtual circuit created in Step 5 to the virtual switch created in Step 1. <pre>virtual-switch ethernet attach mpls-vc <name> vs <name></pre> —end—

Example

The following sample configuration file segment shows the VPLS configuration for the 5150 named R3 shown in [Figure 14-2 on page 14-5](#).

```

!-----vpls-100
virtual-switch ethernet create vs vpls-100 mode vpls
virtual-switch ethernet add vs vpls-100 port 8 vlan 100

!-----Peer: 2.2.2.2
! ---The path is associated with the tunnel. The purpose of the path is
to allow configuration or strict routing for the LSP.
rsvp-te path create rsvp-path R3_R2.path
rsvp-te path set rsvp-path R3_R2.path index 5 ip 10.103.34.4 hop-type loose

! ---Record route information is used to calculate Fast Re-route paths
and therefore should always be enabled if FRR protection is desired.
mpls tunnel create rsvp-ingress R3_R2 dest-ip 2.2.2.2 explicit-tunnel-path
R3_R2.path record-route on

! --- The pw-mode MESH is what allows the ELAN behavior associated
with a VPLS. The pw-type ETH-RAW along with the attachment circuit
type defines the treatment of the outer VLAN of the ingress frame.
mpls l2-vpn create dynamic-vc vpls-100.R2 pw-id 100 peer 2.2.2.2 tunnel R3_R2
pw-mode mesh pw-type eth-raw

virtual-switch ethernet attach mpls-vc vpls-100.R2 vs vpls-100

!-----Peer: 5.5.5.5
! ---Path to peer
rsvp-te path create rsvp-path R3_R5.path
rsvp-te path set rsvp-path R3_R5.path index 5 ip 10.103.34.4 hop-type loose

! ---MPLS Tunnel
mpls tunnel create rsvp-ingress R3_R5 dest-ip 5.5.5.5 explicit-tunnel-path
R3_R5.path record-route on

! --- Virtual Circuit
mpls l2-vpn create dynamic-vc vpls-100.R5 pw-id 100 peer 5.5.5.5 tunnel R3_R5
pw-mode mesh pw-type eth-raw
virtual-switch ethernet attach mpls-vc vpls-100.R5 vs vpls-100

!-----Peer: 6.6.6.6
! ---Path to peer
rsvp-te path create rsvp-path R3_R6.path
rsvp-te path set rsvp-path R3_R6.path index 5 ip 10.103.34.4 hop-type loose

! ---MPLS Tunnel
mpls tunnel create rsvp-ingress R3_R6 dest-ip 6.6.6.6 explicit-tunnel-path
R3_R6.path record-route on

! --- Virtual Circuit
mpls l2-vpn create dynamic-vc vpls-100.R6 pw-id 100 peer 6.6.6.6 tunnel R3_R6
pw-mode mesh pw-type eth-raw

virtual-switch ethernet attach mpls-vc vpls-100.R6 vs vpls-100

```

Procedure 14-4

Configuring a 39XX/51XX H-VPLS

A “[Hierarchical VPLS](#)” on [page 2-7](#) (H-VPLS) model is used to allow spoke connections to the VPLS core. A device provides the functionality to interface with the VPLS core by functioning as an MTU-s, which is connected as a spoke in the VPLS core using a virtual circuit.

An MTU-s connects to an PE on the NNI side through an MPLS network and to a CE on the UNI side through an untagged interface or an 802.1Q interface.


Note: The system handles the outer VLAN which could be C-VLAN, S-VLAN or B-VLAN and not the combination of multiple VLANs as in Q-in-Q or MAC-in-MAC. This is true for a UNI port on a PE also.

Although MTU-s is an external device to a PE, it is viewed as an adjunct device to the PE and used to address scalability.

Note: This procedure assumes that the device has already been configured with base LSR functionality as illustrated in “[Configuring a 39XX/51XX LSR](#)” on [page 14-2](#).

Figure 14-4 illustrates the H-VPLS topology which includes virtual circuit redundancy for protection switching.

Figure 14-4
H-VPLS topology



Step	Action
1	Create a virtual switch. <pre>virtual-switch ethernet create vs <vs-name> mode <vpws vpws></pre>
2	Attach the UNI port(s) to the virtual switch. <pre>virtual-switch ethernet add vs <vs-name> port <name> [vlan <vlan-id>]</pre>
3	Create an explicit path for each endpoint. <pre>rsvp-te path create rsvp-path <name> rsvp-te path set rsvp-path <name> index <#> ip <ip> hop- type <loose strict></pre>
4	Create an MPLS tunnel for each endpoint. <pre>mpls tunnel create rsvp-ingress <name> dest-ip <ip> explicit-tunnel-path <name> record-route <on off></pre>
5	Create a primary and protection virtual circuit (Pseudowire). <pre>mpls l2-vpn create dynamic-vc <vc-name> pw-id <value> peer <peer-ip-addr> tunnel <tunnel-name></pre>
6	Attach the each virtual circuit to the virtual switch created in Step 1. <pre>virtual-switch ethernet attach vs <vs-name> mpls-vc <vc- name></pre>

—end—

Example

The following sample configuration file segment shows the H-VPLS configuration for the 3930 named R1 shown in [Figure 14-4 on page 14-14](#).

```

H-VPLS Configuration Example on R1
!-----hvpls-200
!----Since VPLS is a multipoint service the virtual switch mode should
be set to VPLS which will allow multiple attachment circuits. VPWS
mode only allows a single attachment circuit.
virtual-switch ethernet create vs hvpls-200 mode vpls
virtual-switch ethernet add vs hvpls-200 port 8 vlan 200

!-----Primary Virtual Circuit: hvpls-200P
! --- Primary Path
rsvp-te path create rsvp-path R1_R2.path
rsvp-te path set rsvp-path R1_R2.path index 5 ip 10.101.12.2 hop-type loose

!-----Primary Tunnel
mpls tunnel create rsvp-ingress R1_R2 dest-ip 2.2.2.2 explicit-tunnel-path
R1_R2.path record-route on

! ---- Virtual Circuit, Virtual Switch, and Attachment Circuit
mpls l2-vpn create dynamic-vc hvpls-200P pw-id 200 peer 2.2.2.2 tunnel R1_R2
pw-mode spoke pw-type eth-raw

virtual-switch ethernet attach mpls-vc hvpls-200P vs hvpls-200

! ---- Backup Path
rsvp-te path create rsvp-path R1_R3.path
rsvp-te path set rsvp-path R1_R3.path index 5 ip 10.104.13.3 hop-type loose

!-----Backup Tunnel
mpls tunnel create rsvp-ingress R1_R3 dest-ip 3.3.3.3 explicit-tunnel-path
R1_R3.path record-route on

! ---- Virtual Circuit, Virtual Switch, and Attachment Circuit
mpls l2-vpn protection create dynamic-vc hvpls-200B primary-vc-name hvpls-
200P secondary-pw-id 200 peer 3.3.3.3 tunnel R1_R3 virtual-switch ethernet
attach mpls-vc hvpls-200B vs hvpls-200

```

Interface configuration

This chapter provides the following procedures:

- “Creating an MPLS remote management interface” on page 15-2
- “Displaying interface information” on page 15-5
- “Deleting a management virtual switch” on page 15-7
- “Adding a static ARP entry” on page 15-8
- “Displaying static ARP entries” on page 15-9
- “Allocating resources for an MPLS management virtual switch (3916, 3926m, 3928, 3930 and 3931 platforms)” on page 15-10

Procedure 15-1

Creating an MPLS remote management interface

Create an MPLS remote management interface to provide an in-band communication link to the 39XX/51XX Service Delivery, Aggregation and Virtualization Switch that is accessible from a remote location.

Note 1: There can be only one in-band remote management interface defined at any given time. Therefore, if the CLI session you are using is connected to a remote interface, you will lose access to the current session as the new remote interface is configured, and you will have to initiate a new session over the newly configured virtual switch.

Note 2: In order to carry remote interface traffic the remote-interface is associated with a virtual switch. Management access to the switch can then be gained from any of the members of this virtual switch, including attachment circuit members. Thus, if attachment circuits exist on the virtual switch that is associated with the remote interface, customers could obtain management access to the node.

To prevent this, create a virtual switch specifically for use for in-band management, and do not attach customer attachment circuits to that virtual switch.

Step	Action
1	Create a virtual switch: <pre>virtual-switch create vs <vs></pre>
2	Create the MPLS virtual circuit. <pre>mpls l2-vpn create static-vc <matching-vc></pre>
3	Attach the MPLS virtual circuit. <pre>virtual-switch interface attach mpls-vc <mpls-vc> vs <vs></pre> where mpls-vc is the name of the MPLS virtual circuit <mpls-vc> vs <vs> is the name of the virtual switch
4	Create a remote management interface: <pre>interface remote set ip <ip address> interface remote set vs <vs-name></pre> Note: The virtual switch is of type VPLS.

- 5 The remote is now an untagged service. The following configuration is required on the attachment circuit port:

```
virtual-switch ethernet add vs <vs-name> port <logical-  
port-name> vlan <vlan-id> ingress-l2-transform pop  
egress-l2-transform push-8100.10
```

where

vs <vs-name> is the virtual switch name.

port <logical-port-name> is the logical port name.

vlan <vlan-id> is the VLAN ID.

Note: The port and VLAN should belong to the same port.

—end—

Procedure 15-2

Displaying remote interface configuration

Display the remote interface configuration to identify the Management Domain used for connectivity, which is either a VLAN or a virtual switch.

Step	Action

1	Display information about the remote interface:

```
interface remote show
```

—end—

Example

In the example output below, the remote interface connectivity is provided by an MPLS management virtual switch named "MyMngmt1"

INTERFACE STATE				
Parameter	Value	Source	State	
Name	remote			
Index	15			
Admin State	Enabled			
Oper State	Enabled			
MAC Address	00:02:5a:01:c5:4f			
Management Domain	VS MyMngmt1			
Priority	7			
MTU	1500			
IPv4 Oper addr/mask	192.168.50.1/24	Manual	PREFERRED	
IPv4 Broadcast Address	192.168.50.255	Manual	PREFERRED	
IPv6 Oper addr/mask	fe80::202:5aff:fe01:c54f/64	Internal	PREFERRED	

Procedure 15-3

Displaying interface information

Display interface information to verify interface configuration.

Step	Action

To display interface information

- 1 Display information about interfaces:
`interface show {gateway}`
where
gateway shows gateway information

To display information about a specific interface

- 2 Display information about a specific interface:
`interface show interface <interface>`
where
interface is the interface name
<interface>

To display information about a specific ip-interface

- 3 Display information about an ip-interface:
`interface show ip-interface <ip-interface>`
where
ip-interface <ip-interface> is the ip-interface name
interface>

—end—

Example

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

INTERFACE MANAGEMENT		
Name	Type	IP Address/Prefix
iom.dcn2	partnr	172.16.233.216/24
iom.dcn2	partnr	fe80::2e39:c1ff:fe49:fc03/64
active	active	10.18.2.27/24
active	active	fe80::2e39:c1ff:fe49:fc01/64

TCP/IP/STACK OPERATIONAL STATE	
Parameter	Value
IPv4 Gateway	10.18.2.1
IP Forwarding	on
DCN Auto Revert	off

IPV6 STACK STATE	
Parameter	Value
IPv6 Stack	Enabled
Stack Preference	IPv6
Accept Router Advertisement	On
ICMP Accept Redirects	Off
ICMP Echo Ignore Broadcasts	On
ICMP Port Unreachable	On
Max SLAAC Addresses	16

L3 INTERFACE OPERATIONAL STATE				
Name	Type	IP Address/Prefix	Admin State	Oper State
lbnIf	loop	6.6.6.6/32	Enabled	Enabled
Port21Vlan1800	Ether	67.14.0.162/24	Enabled	Enabled
Port21Vlan1810	Ether	67.14.10.162/24	Enabled	Enabled
Port22Vlan1900	Ether	68.14.0.2/24	Enabled	Enabled
Port22Vlan1910	Ether	68.14.10.2/24	Enabled	Enabled
saos_loop1	loop	172.30.73.252/32	Enabled	Enabled

SERIAL PORT STATE			
Name	Value		
	Primary	Secondary	
Console	Enabled	Enabled	
Debug	Enabled	Enabled	

Procedure 15-4

Deleting a management virtual switch

A virtual switch that is currently associated with the MPLS remote interface cannot be deleted until the remote interface is removed.

Step	Action
1	Delete the remote interface: <pre>interface delete interface <remote-name></pre> where <remote-name> is the name of the remote interface
2	Detach all member MPLS VC interfaces: <pre>virtual-switch interface detach mpls-vc <vc-name></pre> where <vc-name> is the name of the MPLS VC
3	Detach the CPU interface from the management VS: <pre>virtual-switch interface detach cpu-subinterface F<CpuSubInterface></pre> where <vc-name> is the name of the MPLS VC <CpuSubInterface> is the name of the CPU interface
4	Delete the management virtual switch (VS). <pre>virtual-switch delete vs <vs_name></pre> where <vs-name> is the name of the management VS

—end—

Procedure 15-5

Adding a static ARP entry

ARP maps an IP address to a physical MAC address. Entries can be learned dynamically, such as when a device sends an ARP request or ping, but for the entries to be permanent, as in the case of using a device as a static reflector for RFC 2544, you can add a static entry.

Step	Action
1	Add a static ARP entry:<pre>arp static add {destination <IpAddress>} {mac <MacAddress>} {ip-interface <Interface>}</pre> where destination <IpAddress> is the destination host IP address. mac <MacAddress> is the destination MAC address. ip-interface <Interface> is the IP interface routed to the destination. —end—

Procedure 15-6

Displaying static ARP entries

Display static ARP entries.

Step	Action

1	Display static ARP entries:

```
arp static show
```

—end—

Example

The following example shows sample output for the arp static show command.

```
> arp static show
```

DestinationIp	DMAC	ifName (ifIndex)
1.1.1.2	00:02:5a:01:b3:c6	Ref_Test (5)

Procedure 15-7

Allocating resources for an MPLS management virtual switch (3916, 3926m, 3928, 3930 and 3931 platforms)

On the 3916, 3926m, 3928, 3930 and 3931 platforms, support for an MPLS management virtual switch requires adjustment of the default resource allocation:

- If an attempt is made to use the “interface remote set vs <vs-name>” command before these resources have been allocated on the 3916/26m/28/30/31 platforms, an error message is displayed and the command is refused.
- If an attempt is made to remove or reduce the resources allocated to the transport-oam feature set while the remote interface is associated with a virtual switch, the command is refused.

Step	Action
1	Allocate resources to the transport-oam feature set: <pre>resource-manager pool set feature transport-oam resource classifier count 256</pre> <pre>resource-manager pool set feature transport-oam resource counter count 256</pre>
2	Deallocate resources from another feature (do one of the following two options): • Reduce resources from the traffic-profiling feature: <pre>resource-manager pool set feature traffic-profiling resource classifier count 768</pre> <pre>resource-manager pool set feature traffic-profiling resource counter count 768</pre> • Eliminate resources from either the broadcast-containment or vc-statistics feature: <pre>resource-manager pool set feature <feature-name> resource classifier count 0</pre> <pre>resource-manager pool set feature <feature-name> resource counter count 0</pre> feature <feature-name> either “broadcast-containment” or “vc-statistics”
3	Save the configuration and reboot: <pre>config save</pre> <pre>reboot now</pre>

—end—

MPLS troubleshooting

This chapter explains how to troubleshoot MPLS by means of

- forwarding information base (FIB), which stores static route and forwarding information
- adjacency information base (AIB), which stores ARP static route and adjacency information
- traceroute
- ping

This chapter provides the following procedures:

- [“Displaying FIB information” on page 16-2](#)
- [“Displaying the AIB table” on page 16-3](#)
- [“Displaying FIB entries” on page 16-4](#)
- [“Clearing all FIB or AIB entries” on page 16-6](#)
- [“Enabling logging of FIB or AIB events” on page 16-7](#)
- [“Modifying the AIB timeout for MPLS adjacency” on page 16-8](#)
- [“Running a traceroute on a tunnel” on page 16-9](#)
- [“Running a traceroute on a virtual circuit” on page 16-14](#)
- [“Running a VCCV traceroute” on page 16-17](#)
- [“Running ping on a tunnel” on page 16-20](#)
- [“Running ping on a virtual circuit” on page 16-26](#)
- [“Running VCCV ping for MPLS virtual circuits” on page 16-28](#)
- [“Switching over to the backup TE tunnel” on page 16-31](#)
- [“Switching over to the backup TP tunnel” on page 16-32](#)
- [“Manually forcing a switchover to the standby pseudowire and reverting it back” on page 16-33](#)

Procedure 16-1

Displaying FIB information

Display FIB information for troubleshooting purposes. You can display the following information for the FIB:

- log
- interfaces
- interface IP addresses
- static resolutions

Step	Action

To display the FIB log

- 1 Display the FIB log:

```
ip fib log show
```

To display the FIB interfaces

- 2 Display the FIB interfaces:

```
ip fib interfaces show [summary] [details]
```

where

summary	displays a summary of FIB information for interfaces.
details	displays detailed FIB information for interfaces.

To display the FIB interface IP addresses

- 3 Display the FIB interface IP addresses:

```
ip fib interfaceIpAddresses show [summary] | [details]
```

where

summary	displays a summary of FIB information for interface IP addresses.
details	displays detailed FIB information for interface IP addresses.

To display the FIB static resolutions

- 4 Display the FIB static resolutions:

```
ip fib staticResolutions show [summary] | [details]
```

where

summary	displays a summary of FIB information for static resolutions.
details	displays detailed FIB information for static resolutions.

—end—

Procedure 16-2 Displaying the AIB table

Display the Adjacency Information Base (AIB) table to view adjacencies for tunnels and AIS when the next hop is reachable.

Step	Action
1	Display the AIB: <pre>ip aib show [summary] [details]</pre> where summary displays a summary of AIB information. details displays detailed AIB information. —end—

Example

The following example shows sample output for the ip aib show command.

```
> ip aib show
+-----+-----+-----+-----+
| NexthopIp | ifIndex | Dmac | ObjIdx |
+-----+-----+-----+-----+
| 1.1.1.2 | 5 | 00:02:5a:01:b3:c6 | 4 |
+-----+-----+-----+-----+
```

Procedure 16-3

Displaying FIB entries

After creating a static IP route, the forwarding information is stored in the FIB. The summary view is displayed by default.

You can display:

- the FIB entry for the default route
- a summary of FIB entries
- details for each FIB entry

Step	Action

To display the FIB entry for the default route

- 1 Display the FIB entry for the default route:

```
ip fib show [destination <Ip Address with mask>] [default]
```

where

destination is the destination route IP address/mask.
<Ip Address with mask>

[default] displays the default route, if any configured.

To display a summary of FIB entries

- 2 Display a summary of FIB entries:

```
ip fib show [destination <Ip Address with mask>] [summary]
```

where

destination is the destination route IP address/mask.
<Ip Address with mask>

[summary] displays summary.

To display details for each FIB entry

- 3 Display details for each FIB entry:

```
ip fib show [destination <Ip Address with Mask>] [details]  
[summary]
```

where

destination is the destination route IP address.

<Ip Address with
Mask>

[details] displays details.

—end—

Procedure 16-4

Clearing all FIB or AIB entries

Clear all FIB or AIB entries as part of a clean up operation, or when troubleshooting.

Step	Action

To clear all FIB entries

1	Clear all FIB entries: <code>ip fib flush</code>

To clear all AIB entries

2	Clear all AIB entries: <code>ip aib flush</code>

—end—

Procedure 16-5

Enabling logging of FIB or AIB events

Enable logging of FIB or AIB events for troubleshooting.

Step	Action

To enable logging of FIB events

1	Enable logging of FIB events: <code>ip fib debug events on</code>

To enable logging of AIB events

2	Enable logging of AIB events: <code>ip aib debug events on</code>

—end—

Procedure 16-6

Modifying the AIB timeout for MPLS adjacency

Modify the AIB timeout for MPLS adjacency when an AIB timeout longer than the default value is required for the network configuration. An AIB timeout longer than the default value may be required when a broadcast storm occurs in the network. The default value is 60 seconds.

Step	Action
1	Modify the AIB timeout: <pre>ip aib set timeout <SECONDS: 60-600></pre> —end—

Procedure 16-7

Running a traceroute on a tunnel

Run traceroute to test MPLS tunnels.

You can run a traceroute on:

- ingress MPLS-TE tunnels
- ingress MPLS-TP bidirectional tunnels
- ingress MPLS-TP unidirectional tunnels
- egress MPLS-TP bidirectional tunnels
- MPLS-TE bidirectional associated tunnels
- MPLS-TP bidirectional associated tunnels

Table 16-1 shows the default reply modes along with default encapsulation for MPLS traceroute over different types of tunnels.

Table 16-1
Default reply modes and encapsulation for MPLS traceroute

	Tunnel type			
	unidirectional	TE associated	TP associated	TP co-routed
Reply mode	IPv4	IPv4	IPv4	LSP
Encapsulation	IP/UDP	IP/UDP	IP/UDP	IP/UDP

Table 16-2 lists attributes for MPLS traceroute commands.

Table 16-2
Attributes for MPLS traceroute commands

Attribute	Description	Default value
timeout <NUMBER: 500-10000>]	is the timeout value. Valid values are in the range of 500 to 10000 milliseconds.	1000
ttl <NUMBER: 1-30>	is the time to live. Valid values are numbers in the range of 1 to 30.	30
encap <IP/UDPI Non-IP/UDP>	is the encapsulation type.	IP/UDP
validate-fec-stack <on/off>	indicates whether the FEC is validated in the request	on
validate-reverse-fec-stack <on/off>	indicates whether the FEC is validated in the response.	on

Table 16-3 lists output strings.

Table 16-3
Output strings

Output string	Description
! Seq: <number> Latency: <number> ms	Displays '!' with Latency if the ping is successful.
! Loc Hop: 0 NextHopIP: <ip address> NextHopLabel: <Label> MTU: <mtu value> bytes	Displayed if the MPLS traceroute operation is successful.
! LSR Hop: <number> Next Hop: ip <ip address unknown> Label: <Label> MTU: <mtu value> bytes Latency: <number> ms	Displayed if the MPLS ip/udp encapsulated traceroute operation of the LSR is successful. If the IP address of the LSR is unknown, "unknown" is displayed.
! LSR Node: <number> Id: <loopback address> Next Hop: ip <ip address unknown> Label: <Label> MTU: <mtu value> bytes Latency: <number> ms	Displayed when an MPLS non-ip/udp encapsulated traceroute output of the LSR is successful. If the IP address of the LSR is unknown, "unknown" is displayed.
! LSR Hop: <number> Next Hop: ip <ip address unknown> Label: 3 (Implicit NULL) MTU: <mtu value> bytes Latency: <number> ms	Displayed if ip/udp encapsulated traceroute reply is successfully received from a PHP tunnel. If IP address is unknown, "unknown" is displayed.

Table 16-3
Output strings

Output string	Description
! LSR Node: <number> Id: <loopback address> Next Hop: Label: 3 (Implicit NULL) MTU: <mtu value> bytes Latency: <number> ms	Displayed if non-ip/udp traceroute reply is successfully received from a PHP tunnel.
! Egr Hop: <number> Latency: <number> ms	Displayed when an ip/udp encapsulated traceroute reply is received from an egress node.
! Egr Node: <number> Id: <ip address> Latency: <number> ms	Displayed if a non-ip/udp encapsulated traceroute reply is successfully received from an egress node.
! Loc Node: 0 Id <loopback address> Next Hop: Label: <Label> MTU <mtu value> bytes	Displayed in a MS-PW when the non-ip traceroute output of a local node is successful.
! S-PE Node: <number> ip <ip address> unknown> Next S-PE: Label <Label> MTU <mtu value> bytes Latency: <number> ms	Displayed in a MS-PW when the ip/udp encapsulated traceroute output from an S-PE node is successful.
! S-PE Node: <number> Id <loopback address> Next S-PE: Label <Label> MTU <mtu value> bytes Latency: <number> ms	Displayed in a MS-PW when a non-ip/udp traceroute output from the S-PE node is successful.
! Remote T-PE: <number> Latency: <number> ms	Displayed in a MS-PW when ip/udp encapsulated traceroute output from the T-PE node is successful.
! Remote T-PE: <number> Node Id: <loopback address> Latency: <number> ms	Displayed in a MS-PW when a non/ip/udp encapsulated traceroute output from the T-PE node is successful.
X Seq: <number> Received an undetermined error	Displayed when a return code of 0 is received by the requester.
X Seq: <number> Malformed echo request frame	Returned by the receiver when the receiver is not able to understand the MPLS Echo payload.
X Seq: <number> One or more of the TLVs was not understood	Returned by the receiver when the receiver is not able to understand one or more TLVs.

Table 16-3
Output strings

Output string	Description
X Seq: <number> Replying router has no mapping for the FEC	Returned by the receiver when the receiver is not able to find a match for a FEC specified in the MPLS Echo payload. If this error is received, validate the configuration on both the LERs.
XSeq <number> Downstream Mapping Mismatch	Returned by the received when the DSMAP in the MPLS Echo request package does not match with the receiver's information. This is only returned as part of MPLS traceroute.
X Seq: <number> TTL expired at Label Switched Router	Displayed when the replying router is an LSR for the MPLS ping request. This is only returned as part of MPLS Ping.
XSeq: <number> TTL expired at PHP	Displayed when the replying router is a PHP router. This is only returned as part of MPLS Ping.
X Seq: <number> Mapping for FEC is not given label	Returned by the receiver when the receiver is able to find a match for the FEC but labels corresponding to the FEC do not match.
X Seq: <number> No Label entry	Returned by the receiver when the receiver does not have an entry for the label received.
X Seq: <number> Protocol not associated with interface	Returned by the receiver when the receiver is not able to find a sub-TLV with respect to an interface.
X Seq: <number> Premature termination of ping due to label stack shrinking to a single label	Displayed when there is a premature termination of the ping operation due to label stack shrinking to a single label.
X Seq <number> Request Timed out	Displayed when a reply is not received in the timeout amount of time.

Step	Action

To run a traceroute on an ingress MPLS-TE tunnel

- 1 Run a traceroute on an ingress MPLS-TE tunnel:

```
mpls traceroute tunnel <tunnel> [timeout <NUMBER: 500-10000>] [ttl <NUMBER: 1-30>] [validate-fec-stack <on|off>] [validate-reverse-fec-stack <on|off>]
```

To run a traceroute on an ingress MPLS-TP bidirectional tunnel

- 2 Run a traceroute on an ingress MPLS-TP bidirectional tunnel:

```
mpls traceroute tp-tunnel-ingress-corout <tp-tunnel-ingress-corout> [timeout <NUMBER: 500-10000>] [ttl <NUMBER: 1-30>] [reply-mode <ipv4|lsp>] [encap <ip/udp|non-ip/udp>] [validate-fec-stack <on|off>] [validate-reverse-fec-stack <on|off>]
```

To run a traceroute on an ingress MPLS-TP unidirectional tunnel

- 3 Run a traceroute on an ingress MPLS-TP unidirectional tunnel:

```
mpls traceroute tp-tunnel-ingress-unidir <tp-tunnel-ingress-unidir> [timeout <NUMBER: 500-10000>] [ttl <NUMBER: 1-30>] [validate-fec-stack <on|off>] [validate-reverse-fec-stack <on|off>]
```

To run a traceroute on an egress MPLS-TP bidirectional tunnel

- 4 Run a traceroute on an egress MPLS-TP bidirectional tunnel:

```
mpls traceroute tp-tunnel-egress-corout <tp-tunnel-egress-corout> [timeout <NUMBER: 500-10000>] [ttl <NUMBER: 1-30>] [reply-mode <ipv4|lsp>] [encap <ip/udp|non-ip/udp>] [validate-fec-stack <on|off>] [validate-reverse-fec-stack <on|off>]
```

To run a traceroute on an MPLS-TE bidirectional associated tunnel

- 5 Run a traceroute on an MPLS-TE bidirectional associated tunnel:

```
mpls traceroute assoc-lsp <assoc-lsp> [timeout <NUMBER: 500-10000>] [ttl <NUMBER: 1-30>] [validate-fec-stack <on|off>] [validate-reverse-fec-stack <on|off>]
```

To run a traceroute on an MPLS-TP bidirectional associated tunnel

- 6 Run a traceroute on an MPLS-TP bidirectional associated tunnel:

```
mpls traceroute assoc-tp-lsp <assoc-tp-lsp> [timeout <NUMBER: 500-10000>] [ttl <NUMBER: 1-30>] [validate-fec-stack <on|off>] [validate-reverse-fec-stack <on|off>]
```

—end—

Procedure 16-8

Running a traceroute on a virtual circuit

Run a traceroute to test a virtual circuit.

[Table 16-4](#) lists output strings.

Table 16-4
Output strings

Output string	Description
! Seq: <number> Latency: <number> ms	Displays '!' with Latency if the ping is successful.
! Loc Hop: 0 NextHopIP: <ip address> NextHopLabel: <Label> MTU: <mtu value> bytes	Displayed if the MPLS traceroute operation is successful.
! LSR Hop: <number> NextHopIP: <ip address> NextHopLabel: <Label> MTU: <mtu value> bytes	Displayed if the MPLS traceroute operation is successful.
! LSR Hop: <number> NextHopIP: <ip address> NextHopLabel: 3 (Implicit NULL) MTU: <mtu value> bytes	Displayed when a traceroute reply is received from a PHP tunnel.
! Egr Hop: <number> Latency	Displayed when a traceroute reply is received from an egress tunnel.
X Seq: <number> Received an undetermined error	Displayed when a return code of 0 is received by the requester.
X Seq: <number> Malformed echo request frame	Returned by the receiver when the receiver is not able to understand the MPLS Echo payload.
X Seq: <number> One or more of the TLVs was not understood	Returned by the receiver when the receiver is not able to understand one or more TLVs.
X Seq: <number> Replying router has no mapping for the FEC	Returned by the receiver when the receiver is not able to find a match for a FEC specified in the MPLS Echo payload. When this error is received, validate the configuration on both the LERs.
X Seq: <number> Downstream Mapping Mismatch	Returned by the receiver when the DSMAP in the MPLS Echo request packet does not match the information the receiver has.
X Seq: <number> TTL expired at Label Switched Router	Displayed when the replying router is an LSR for the MPLS ping request.
X Seq: <number> TTL expired at PHP	Displayed when the replying router is a PHP router.

Output string	Description
X Seq: <number> Mapping for FEC is not given label	Returned by the receiver when the receiver is able to find a match for the FEC but labels corresponding to the FEC do not match.
X Seq: <number> No Label entry	Returned by receiver when the receiver does not have an entry for the label received.
X Seq: <number> Protocol not associated with interface	Returned by the receiver when the receiver is not able to find a sub-TLV with respect to an interface.
X Seq: <number> Premature termination of ping due to label stack shrinking to a single label	Displayed when there is a premature termination of the ping operation due to label stack shrinking to a single label.
X Seq: <number> Request Timed out	Displayed when a reply is not received in the timeout amount of time.

Step	Action

1 Run a traceroute on a virtual circuit:

```
mpls traceroute vc <vc> [timeout <NUMBER: 500-10000>]
[ttl <NUMBER: 1-30>] [encap <ip/udp | non-ip/udp>]
[validate-fec-stack <on | off>] [validate-reverse-fec-
stack <on | off>]
```

where

vc <vc> is the name of the MPLS virtual circuit

timeout is the timeout in milliseconds

<MILLISECONDS:
500-10000>

ttl <NUMBER: 1-30> is the time-to-live

encap <ip/udp | non-ip/udp> is the type of encapsulation

validate-fec-stack <on | off> indicates whether the FEC is validated in the request

validate-reverse-fec-stack <on | off> indicates whether the FEC is validated in the response

—end—

Procedure 16-9

Running a VCCV traceroute

Run a VCCV traceroute on a SS-PW or a MS-PW.

Step	Action

To run a traceroute on an MPLS SS-PW

- Run a traceroute on a SS-PW:

```
mpls traceroute vc <vc-name> [timeout <NUMBER: 500-10000>] [ttl <NUMBER: 1-30>]
```

where

vc <vc-name> is the MPLS pseudowire name

timeout is the timeout in milliseconds
<NUMBER: 500-10000>

ttl <NUMBER: 1-30> is the time-to-live

To run a traceroute on an MPLS MS-PW

- Run a traceroute on a MS-PW:

```
mpls traceroute vc <vc-name> [timeout <NUMBER: 500-10000>] [ttl <NUMBER: 1-255>] {segment <num-of-segment-hops>} [fec-stack-validation <on | off>] [reverse-fec-stack-validation <on | off>] [dsmap-validation <on | off>]
```

where

vc <vc-name> is the MPLS pseudowire name

timeout is the timeout in milliseconds
<NUMBER: 500-10000>

ttl <NUMBER: 1-255> is the time-to-live

segment <num-of-segment-hops> is the last segment number

where

fec-stack-validation <on off>	enable or disable FEC validation
reverse-fec-stack-validation <on off>	enable or disable FEC validation in the reverse path
dsmap-validation <on off>	enable or disable downstreamMap validation

To validate a traceroute on an MPLS MS-PW

- 3** Validate a traceroute on an MPLS-TP MS-PW:
- ```
mpls traceroute vc <vc-name> segment <num-of-segment-hops> [fec-stack-validation <on|off>] [timeout <NUMBER: 500-10000>]
```

where

|                                 |                                  |
|---------------------------------|----------------------------------|
| vc <vc-name>                    | is the MPLS pseudowire name      |
| segment <num-of-segment-hops>   | is the last segment number       |
| fec-stack-validation <on   off> | enable or disable FEC validation |
| timeout <NUMBER: 500-10000>     | is the timeout in milliseconds   |

—end—

## Examples

The following example performs a traceroute on a virtual circuit named vc-to-boston with ip/udp encapsulation.

```
mpls traceroute vc vc-to-boston segment 3 encap ip/udp validate fec stack on
```

Output for the command is:

```
MPLS traceroute with ttl (30) timeout (1000 ms)
Legend: `!' - Success, `X' - Error
! Loc Node: 0 Next S-PE: ip 30.30.30.30 Label 8000 MTU 1500
! S-PE: 1 Next S-PE: ip 40.40.40.40 Label 5000, MTU 1500 Latency: 6 ms
! S-PE: 2 Next S-PE: ip 50.50.50.50 Label 6000, MTU 1500 Latency: 6 ms
! Remote T-PE: 3 Latency: 3 ms
```

The following example performs a traceroute on a virtual circuit named vc-to-boston with non-ip/udp encapsulation.

```
mpls traceroute vc vc-to-boston segment 3 encap non-ip/udp validate-fec-stack
off
```

Output for the command is:

```
MPLS traceroute with ttl (30) timeout (1000 ms
Legend: `!' - Success, `X' - Error
! Loc Node 0: Id 20.20.20.20 Next S-PE: Label 8000 MTU 1500
! S-PE Node: 1 Id 30.30.30.30 Next S-PE: Label 5000 MTU 1500 Latency: 5 ms
! S-PE Node 2: Id 40.40.40.40 Next S-PE: Label 6000 MTU 1500 Latency: 7ms
! Remote T-PE 3: Node Id: 50.50.50.50 Latency: 8 ms

-----Statistics-----
5 packets transmitted, 5 packets received
round-trip (ms) min/avg/max=4/4/5
```

## Procedure 16-10

### Running ping on a tunnel

Run ping to test tunnels. You can ping:

- ingress MPLS-TE tunnels
- ingress MPLS-TP bidirectional tunnels
- ingress MPLS-TP unidirectional tunnels
- egress MPLS-TP bidirectional tunnels
- MPLS-TE bidirectional associated tunnels
- MPLS-TP bidirectional associated tunnels

Table 16-5 shows the default reply modes along with default encapsulation for the ping operation over different types of tunnels.

**Table 16-5**  
**Default reply modes and encapsulation for the ping operation**

|                                                | Tunnel type    |               |               |              |
|------------------------------------------------|----------------|---------------|---------------|--------------|
|                                                | Unidirectional | TE associated | TP associated | TP co-routed |
| Reply mode                                     | IPv4           | LSP           | LSP           | LSP          |
| Encapsulation                                  | IP/UDP         | IP/UDP        | IP/UDP        | IP/UDP       |
| Reply mode when time-to-live (TTL) is provided | IPv4           | IPv4          | IPv4          | LSP          |

Table 16-6 lists attributes for ping commands.

**Table 16-6**  
**Attributes for ping commands**

| Attribute                             | Description                                                                                                                                                 |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| count <NUMBER: 1-100>]                | is the number of packets to send. The default value is 5.                                                                                                   |
| packet-size <NUMBER: 96-1464>]        | is the packet size.<br>Default values are: <ul style="list-style-type: none"> <li>• for a dynamic-tunnel: 96</li> <li>• for a static-tunnel: 118</li> </ul> |
| timeout <NUMBER: 500-10000>           | is the timeout in milliseconds. The default value is 1000.                                                                                                  |
| ttl <NUMBER: 1-255>                   | is the time-to-live. The default value is 255.                                                                                                              |
| validate-fec-stack <on   off>         | indicates whether the FEC is validated in the request.                                                                                                      |
| validate-reverse-fec-stack <on   off> | indicates whether the FEC is validated in the response.                                                                                                     |
| reply-mode <IPv4   LSP>               | is the reply mode. The default value is LSP.                                                                                                                |
| encap <ip/udp   non-ip/udp>           | is the type of encapsulation. The default value is IP/UDP.                                                                                                  |

Table 16-7 lists output strings.

**Table 16-7**  
**Output strings**

| Output string                                                                                      | Description                                                      |
|----------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| ! Seq: <number> Latency: <number> ms                                                               | Displays '!' with Latency if the ping is successful.             |
| ! Loc Hop: 0 NextHopIP: <ip address> NextHopLabel: <Label> MTU: <mtu value> bytes                  | Displayed if the MPLS traceroute operation is successful.        |
| ! LSR Hop: <number> NextHopIP: <ip address> NextHopLabel: <Label> MTU: <mtu value> bytes           | Displayed if the MPLS traceroute operation is successful.        |
| ! LSR Hop: <number> NextHopIP: <ip address> NextHopLabel: 3 (Implicit NULL) MTU: <mtu value> bytes | Displayed when a traceroute reply is received from a PHP tunnel. |



| Output string                                              | Description                                                                                                                                                                                         |
|------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ! Egr Hop: <number> Latency                                | Displayed when a traceroute reply is received from an egress tunnel                                                                                                                                 |
| X Seq: <number> Received an undetermined error             | Displayed when a return code of 0 is received by the requester.                                                                                                                                     |
| X Seq: <number> Malformed echo request frame               | Returned by the receiver when the receiver is not able to understand the MPLS Echo payload.                                                                                                         |
| X Seq: <number> One or more of the TLVs was not understood | Returned by the receiver when the receiver is not able to understand one or more TLVs.                                                                                                              |
| X Seq: <number> Replying router has no mapping for the FEC | Returned by the receiver when the receiver is not able to find a match for a FEC specified in the MPLS Echo payload.<br><br>If this error is received, validate the configuration on both the LERs. |
| X Seq: <number> Upstream Interface Index Unknown           | Displayed when the upstream interface index is not known.                                                                                                                                           |
| X Seq: <number> Reserved/Invalid Return Code               | Displayed when there is a reserved or invalid return code.                                                                                                                                          |
| X Seq: <number> TTL expired at Label Switched Router       | Displayed when the replying router is an LSR for the MPLS ping request.                                                                                                                             |
| X Seq: <number> TTL expired at PHP                         | Displayed when the replying router is a PHP router.                                                                                                                                                 |
| X Seq: <number> Mapping for FEC is not given label         | Returned by the receiver when the receiver is able to find a match for the FEC but labels corresponding to the FEC do not match.                                                                    |
| X Seq: <number> No Label entry                             | Returned by the receiver when the receiver does not have an entry for the label received.                                                                                                           |

| Output string                                                                                | Description                                                                                                           |
|----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| X Seq: <number> Protocol not associated with interface                                       | Returned by the receiver when the receiver is not able to find a sub-TLV with respect to an interface.                |
| X Seq: <number> Premature termination of ping due to label stack shrinking to a single label | Displayed when there is a premature termination of the ping operation due to label stack shrinking to a single label. |
| X Seq: <number> Request Timed out                                                            | Displayed when a reply is not received in the timeout amount of time.                                                 |

| Step | Action |
|------|--------|
|------|--------|

**To ping an ingress MPLS-TE tunnel**

- 1 Ping an ingress MPLS-TE tunnel:  

```
mpls ping tunnel <tunnel> [count <NUMBER: 1-100>]
[packet-size <NUMBER: 96-1464>] [timeout <NUMBER: 500-10000>] [ttl <NUMBER: 1-255>] [validate-fec-stack <on | off>] [validate-reverse-fec-stack <on | off>]
```

**To ping an ingress MPLS-TP bidirectional tunnel**

- 2 Ping an ingress MPLS-TP bidirectional tunnel:  

```
mpls ping tp-tunnel-ingress-corout <tp-tunnel-ingress-corout> [count <NUMBER: 1-100>] [packet-size <NUMBER: 96-1464>] [timeout <NUMBER: 500-10000>] [ttl <NUMBER: 1-255>] [reply-mode <IPv4|LSP>] [encap <IP/UDP|Non-IP/UDP>] [validate-fec-stack <on | off>] [validate-reverse-fec-stack <on | off>]
```

**To ping an ingress MPLS-TP unidirectional tunnel**

- 3 Ping an ingress MPLS-TP unidirectional tunnel:  

```
mpls ping tp-tunnel-ingress-unidir <tp-tunnel-ingress-unidir> [count <NUMBER: 1-100>] [packet-size <NUMBER: 96-1464>] [timeout <NUMBER: 500-10000>] [ttl <NUMBER: 1-255>] [validate-fec-stack <on | off>] [validate-reverse-fec-stack <on | off>]
```

**To ping an egress MPLS-TP bidirectional tunnel**

- 4 Ping an egress MPLS-TP bidirectional tunnel:  

```
mpls ping tp-tunnel-egress-corout <tp-tunnel-egress-corout> [count <NUMBER: 1-100>] [packet-size <NUMBER: 96-1464>] [timeout <NUMBER: 500-10000>] [ttl <NUMBER: 1-255>] [reply-mode <IPv4|LSP>] [encap <IP/UDP|Non-IP/UDP>] [validate-fec-stack <on | off>] [validate-reverse-fec-stack <on | off>]
```

**To ping an MPLS-TE bidirectional associated tunnel**

- 5 Ping an MPLS-TE bidirectional associated tunnel:  

```
mpls ping assoc-lsp <assoc-lsp> [count <NUMBER: 1-100>] [packet-size <NUMBER: 96-1464>] [timeout <NUMBER: 500-10000>] [ttl <NUMBER: 1-255>] [reply-mode <IPv4|LSP>] [validate-fec-stack <on | off>] [validate-reverse-fec-stack <on | off>]
```

---

**To ping an MPLS-TP bidirectional associated tunnel****6** Ping an MPLS-TP bidirectional associated tunnel:

```
mpls ping assoc-tp-lsp <assoc-tp-lsp> [count <NUMBER: 1-100>] [packet-size <NUMBER: 96-1464>] [timeout <NUMBER: 500-10000>] [ttl <NUMBER: 1-255>] [reply-mode <IPv4|LSP>] [validate-fec-stack <on | off>] [validate-reverse-fec-stack <on | off>]
```

—end—

## Procedure 16-11

### Running ping on a virtual circuit

Run ping to test a virtual circuit.

The reply-mode setting in the mpls ping command and the CC type setting for the virtual circuit must agree for the ping operation to proceed. The CC type is specified in the pw-vcv-profile and associated to a virtual circuit when the virtual circuit is created.

If the negotiated CC type is cc-ttl-exp, that is, CC Type-3, set the reply-mode to ipv4. If the CC type is cc-ciena-oob, that is, CC Type-4, set the reply mode to lsp. The CC type must be the same between the two virtual circuit peers. If they are not configured with the same CC type, the ping operation will not proceed.

| Step                        | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |         |                                     |                       |                                                           |                     |                                                        |                             |                                                                  |                     |                                                                            |                       |                    |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|-------------------------------------|-----------------------|-----------------------------------------------------------|---------------------|--------------------------------------------------------|-----------------------------|------------------------------------------------------------------|---------------------|----------------------------------------------------------------------------|-----------------------|--------------------|
| 1                           | <p>Run ping for an MPLS virtual circuit:</p> <pre>mpls ping vc &lt;vc&gt; [count &lt;NUMBER: 1-100&gt;] [packet-size &lt;BYTES&gt;] [timeout &lt;NUMBER: 500-10000&gt;] [fec-128-len &lt;14 16&gt;] [reply-mode &lt;ipv4 lsp&gt;] [encap &lt;ip/udp   non-ip/udp&gt;] [validate-fec-stack &lt;on   off&gt;] [validate-reverse-fec-stack &lt;on   off&gt;]</pre> <p>where</p> <table> <tr> <td>vc &lt;vc&gt;</td><td>is the name of the MPLS pseudowire.</td></tr> <tr> <td>count &lt;NUMBER: 1-100&gt;</td><td>is the number of packets to send. The default value is 5.</td></tr> <tr> <td>packet-size &lt;BYTES&gt;</td><td>is the packet size in bytes. The default value is 138.</td></tr> <tr> <td>timeout &lt;NUMBER: 500-10000&gt;</td><td>is the timeout value in milliseconds. The default value is 1000.</td></tr> <tr> <td>fec-128-len &lt;14 16&gt;</td><td>is the number of bytes to use for FEC 128 length. The default value is 14.</td></tr> <tr> <td>reply-mode &lt;ipv4 lsp&gt;</td><td>is the reply mode.</td></tr> </table> <ul style="list-style-type: none"> <li>• ipv4, for a virtual circuit with CC type set to cc-ttl-exp</li> <li>• lsp, for a virtual circuit with CC type set to cc-ciena-oob</li> </ul> | vc <vc> | is the name of the MPLS pseudowire. | count <NUMBER: 1-100> | is the number of packets to send. The default value is 5. | packet-size <BYTES> | is the packet size in bytes. The default value is 138. | timeout <NUMBER: 500-10000> | is the timeout value in milliseconds. The default value is 1000. | fec-128-len <14 16> | is the number of bytes to use for FEC 128 length. The default value is 14. | reply-mode <ipv4 lsp> | is the reply mode. |
| vc <vc>                     | is the name of the MPLS pseudowire.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |         |                                     |                       |                                                           |                     |                                                        |                             |                                                                  |                     |                                                                            |                       |                    |
| count <NUMBER: 1-100>       | is the number of packets to send. The default value is 5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |         |                                     |                       |                                                           |                     |                                                        |                             |                                                                  |                     |                                                                            |                       |                    |
| packet-size <BYTES>         | is the packet size in bytes. The default value is 138.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |         |                                     |                       |                                                           |                     |                                                        |                             |                                                                  |                     |                                                                            |                       |                    |
| timeout <NUMBER: 500-10000> | is the timeout value in milliseconds. The default value is 1000.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |         |                                     |                       |                                                           |                     |                                                        |                             |                                                                  |                     |                                                                            |                       |                    |
| fec-128-len <14 16>         | is the number of bytes to use for FEC 128 length. The default value is 14.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |         |                                     |                       |                                                           |                     |                                                        |                             |                                                                  |                     |                                                                            |                       |                    |
| reply-mode <ipv4 lsp>       | is the reply mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |         |                                     |                       |                                                           |                     |                                                        |                             |                                                                  |                     |                                                                            |                       |                    |

where

encap <ip/udp |  
non-ip/udp> is the type of encapsulation

validate-fec-stack indicates whether the FEC is validated in the request  
<on | off>

validate-reverse- indicates whether the FEC is validated in the response.  
fec-stack <on |  
off>

—end—

## Procedure 16-12

### Running VCCV ping for MPLS virtual circuits

Run VCCV ping to test MPLS virtual circuits.

The reply-mode setting and vc cc-type setting must agree for vc ping to proceed. The vc cc-type is specified in the pw-vccv-profile and associated to a virtual circuit when the virtual circuit is created.

- If the negotiated cc-type is cc-ttl-exp, the reply-mode must be set to ipv4.
- if the reply-mode is cc-ciena-oob, the reply mode must be lsp.

Also, cc-type must also be the same between the two virtual circuits. If they are not configured with the same cc-type, VCCV ping will not proceed.

| Step | Action |
|------|--------|
|------|--------|

- |   |                                            |
|---|--------------------------------------------|
| 1 | Run VCCV ping for an MPLS virtual circuit: |
|---|--------------------------------------------|

```
mpls ping vc <vc> [count <NUMBER: 1-100>] [packet-size
<NUMBER: 114-1464>] [timeout <NUMBER: 500-10000>] [fec-
128-len <14|16>] [reply-mode <IPv4|LSP>] [encap <ip/
udp|non-ip/udp>] [segment <NUMBER: 1..255>] [seg-src-ip <
IP address>] [seg-dest-ip <IP address>] [seg-pw-id
<NUMBER: 1..4294967295>] [fec-stack-validation
<on|off>] {reverse-fec-stack-validation <on|off>}
```

where

vc <vc> is the MPLS pseudowire name.

count  
<NUMBER: 1-100> is the number of packets to send. The default value is 5.

packet-size  
<NUMBER: 114-1464> is the packet size in bytes. The default value is 114.

timeout  
<NUMBER: 500-10000> is the timeout value in milliseconds. The default value is 1000.

fec-128-len <14|16> is the number of bytes to use for FEC 128 length. The default value is 14.

reply-mode <IPv4|LSP> is the reply mode.

- For TE tunnels, no-op.
- For TP tunnels, the default value is LSP.

encap <ip/udp|non-ip/udp> selects the different types of encapsulation for MS-PW only.

where

segment specifies the target segment to ping to for MS-PW only. The default value is 255.

<NUMBER:

1..255>

The following outlines what occurs in various scenarios:

- If segment, seg-src-ip, seg-dest-ip and seg-pw-id are not set, the segment defaults to 1. The rest of the parameters are filled by SAOS with the information of the first segment.
- If segment is set to a number greater than 1 to ping a MS-PW, but seg-src-ip, seg-dest-ip and seg-pw-id are not set, FEC validation will fail. An error message is displayed.
- If the last segment FEC parameters are specified, but segment is not, VCCV ping will not be initiated. The three FEC parameters must be specified together. VCCV ping will not be initiated. An error message is displayed.

seg-src-ip <IP address>

is the last segment source IP address for MS-PW only.

seg-dest-ip <IP address>

is the last segment destination IP address for MS-PW only.

seg-pw-id  
<NUMBER:  
1..4294967295>

is the last segment PW ID for MS-PW only. The default is LSP.

fec-stack-validation  
<on/off>

enables or disables FEC validation. This parameter specifies if the target node of the ping request packet should or should not verify the FEC content in the target FEC stack TLV.

The following outlines what occurs in various scenarios:

- If fec-stack-validation is on, segment, seg-src-ip, seg-dest-ip and seg-pw-d must be set for MS-PW only.
- If these parameters are not set, the segment parameter defaults to 1. VCCV ping is initiated over the first segment only. The other segment parameters are automatically filled in with available segment information so that FEC validation is correct.

reverse-stack-validation  
<on/off>

enables or disables FEC validation in the reverse. For a MS-PW, reverse path FEC validation is not possible, so it is set off automatically by SAOS, overriding the setting from the CLI. If the verification passes, the return code is 3: "Replying router is an egress at stack depth." If the verification fails, an error, one of the error codes in [Table 16-7](#) is displayed.

—end—

## Examples

The following example configures MS-PW VCCV ping with ip/udp encap and shows the output.



## 16-30 MPLS troubleshooting

---

```
mpls ping vc vc-to-boston segment 2 seg-src-ip 30.30.30.30 seg-dest-ip
40.40.40.40 seg-pw-id 44
```

```
MPLS Ping with count (5) timeout (1000 ms) size (138 bytes)
```

```
Legend: '!' - Success, 'X' - Error
```

```
! Seq: 1 Latency: 13 ms
```

```
! Seq: 2 Latency: 4 ms
```

```
! Seq: 3 Latency: 4 ms
```

```
! Seq: 4 Latency: 4 ms
```

```
! Seq: 5 Latency: 4 ms
```

```
-----Statistics-----
```

```
5 packets transmitted, 5 packets received
```

```
round-trip (ms) min/avg/max=4/5/13
```

The following example configures MS-PW VCCV ping with non-ip/udp encap and shows the output.

```
mpls ping vc si-to-sf segment 2 seg-src-ip 30.30.30.30 seg-dest-ip 50.50.50.50
seg-pw-id 77 encap non-ip/udp
```

```
MPLS Ping with count (5) timeout (1000 ms) size (138 bytes)
```

```
Legend: '!' - Success, 'X' - Error
```

```
! Seq: 1 Latency: 5 ms
```

```
! Seq: 2 Latency: 5 ms
```

```
! Seq: 3 Latency: 4 ms
```

```
! Seq: 4 Latency: 5 ms
```

```
! Seq: 5 Latency: 5 ms
```

```
-----Statistics-----
```

```
5 packets transmitted, 5 packets received
```

```
round-trip (ms) min/avg/max=4/4/5
```

## Procedure 16-13

### Switching over to the backup TE tunnel

You can switch over:

- an MPLS TE-Tunnel
- a static ingress TE-Tunnel
- a static bidirectional ingress associated TE-Tunnel

| Step | Action |
|------|--------|
|------|--------|

#### *To switch over an TE tunnel*

- 1 Switch an active MPLS TE-Tunnel to the backup MPLS TE-Tunnel:  

```
mpls tunnel switchover rsvp-ingress <rsvp-ingress>
```

#### *To switch over a static ingress TE tunnel*

- 2 Switch an active static ingress tunnel to the backup static ingress tunnel:  

```
mpls tunnel switchover rsvp-ingress <rsvp-ingress>
```

#### *To switch over a static bidirectional ingress associated TE Tunnel*

- 3 Switch a static bidirectional ingress associated TE-Tunnel to the backup bidirectional associated TE-Tunnel:  

```
mpls tunnel switchover bidir-ingress-assoc <bidir-ingress-assoc>
```

—end—

## Procedure 16-14

### Switching over to the backup TP tunnel

---

You can switch over a TP tunnel. You can only switch over the active tunnel.

**Note:** If tunnel reversion is on, the system switches from the backup to the primary tunnel once the fault on the primary tunnel is cleared and after waiting the amount of time specified as the reversion hold time (the default is 30 seconds). Do not switch over the backup GMPLS TP tunnel while the reversion hold timer is counting down.

---

| Step | Action |
|------|--------|
|------|--------|

---

- |   |                                                                                                                                   |
|---|-----------------------------------------------------------------------------------------------------------------------------------|
| 1 | Switch a TP tunnel to the backup tunnel:<br><pre>gmpls tp-tunnel switchover rsvp-ingress-unidir &lt;rsvp-ingress-unidir&gt;</pre> |
|---|-----------------------------------------------------------------------------------------------------------------------------------|

—end—

## Procedure 16-15

### Manually forcing a switchover to the standby pseudowire and reverting it back

This procedure allows the operator to gracefully redirect the traffic to the standby pseudowire. This could allow the operator to perform maintenance actions to the primary pseudowire without impacting services.

When manual switchover is turned on, the secondary pseudowire becomes the active pseudowire and a switchover to the secondary pseudowire occurs immediately. If reversion was enabled, it will be disengaged. If the secondary pseudowire is the active pseudowire when the manual switchover is turned on, it simply stays active with the reversion disengaged.

In this mode, though the reversion is disengaged, the pseudowire is still protected. Therefore, if any faults occur on the now active secondary pseudowire, a switchover to the primary pseudowire is performed if the primary is up. Whether the switchover succeeds or not, the manual switchover is canceled as soon as the secondary goes down. The reversion mechanism, if configured for the primary pseudowire, is re-engaged.

When the manual switchover is turned off, the switchover to the primary occurs immediately if the primary is up and the reversion mechanism, if configured for the primary pseudowire is re-engaged.

---

| Step | Action |
|------|--------|
|------|--------|

---

#### ***Manually forcing a switchover to the standby pseudowire***

- 1 Switch over to the dynamic protection virtual circuit:

```
mpls l2-vpn protection switchover <static-vc |
dynamic-vc> <primary_pw_name> on
```

where

<static-vc | dynamic-vc> is the virtual circuit name.  
vc>

<primary\_pw\_name> is the primary PW name.

#### ***Reverting a manual switchover back from the standby pseudowire***

- 2 Switch back from the dynamic protection virtual circuit:

```
mpls l2-vpn protection switchover <static-vc |
dynamic-vc> <primary_pw_name> off
```

—end—

## Procedure 16-16

### Clearing MPLS transit tunnel information

---

You can clear traffic statistics for:

- static transit tunnels
- RSVP transit tunnels

---

| Step | Action |
|------|--------|
|------|--------|

---

#### *To clear traffic statistics for static transit tunnels*

- 1 Clear traffic statistics for static transit tunnels:  

```
mpls tunnel clear static-transit <tunnel name> traffic-
statistics
```

#### *To clear traffic statistics for RSVP transit tunnels*

- 2 Clear traffic statistics for RSVP transit tunnels:  

```
mpls tunnel clear rsvp-transit <tunnel name> statistics
```

—end—

## Procedure 16-17

### Clearing GMPLS TP tunnel information

You can clear traffic statistics for

- static transit tunnels
- RSVP transit tunnels
- static ingress/egress TP co-routed tunnels

| Step | Action |
|------|--------|
|------|--------|

#### **To clear traffic statistics for static TP unidirectional tunnels**

- 1 Clear traffic statistics for static TP tunnels:  

```
gmpls tp-tunnel clear static-transit-undir <tunnel name>
<traffic statistics>
```

#### **To clear traffic statistics for static transit corouted tunnels**

- 2 Clear traffic statistics for static transit tunnels:  

```
gmpls tp-tunnel clear static-transit-corout <tunnel name>
statistics
```

#### **To clear traffic statistics for RSVP transit tunnels**

- 3 Clear traffic statistics for RSVP transit tunnels:  

```
gmpls tunnel clear rsvp-transit-unidir <tunnel name>
<traffic statistics>
```

#### **To clear traffic statistics for static ingress TP co-routed tunnels**

- 4 Clear traffic statistics for static ingress TP co-routed tunnels:  

```
gmpls tp-tunnel clear static-ingress-corout <static-
ingress-corout> statistics
```

#### **To clear traffic statistics for static egress TP co-routed tunnels**

- 5 Clear traffic statistics for static egress TP co-routed tunnels:  

```
gmpls tp-tunnel clear static-egress-corout <static-
egress-corout> statistics
```

—end—

## Procedure 16-18

### Clearing traffic statistics for a virtual circuit

---

You can clear statistics for static and dynamic virtual circuits, including multisegment pseudowire (MS-PW) and the segmented switching provider edge (S-PE) network element (NE) node virtual circuits.

---

| Step | Action |
|------|--------|
|------|--------|

---

- |   |                                           |
|---|-------------------------------------------|
| 1 | Clear virtual circuit traffic statistics: |
|---|-------------------------------------------|

```
mpls l2-vpn clear vc <vc-name> statistics
```

where

vc <vc-name>      is the virtual circuit name.

statistics          clears traffic statistics for static and dynamic virtual circuits

—end—

---

# Bidirectional Forwarding Detection configuration

---

Bidirectional Forwarding Detection (BFD) is a protocol that provides low-overhead, fast detection of link or node failure between any two nodes running over any medium.

The following table shows the BFD session indices in SAOS for 6.x platforms:

**Table 17-1**  
**BFD session indices in SAOS for 6.x platforms**

| Service type       | BFD session type        | Before SAOS 6.18 | After SAOS 6.18 |
|--------------------|-------------------------|------------------|-----------------|
| IPv4 interface     | Single-/Multihop IP-BFD | 1-2256           | 1-256           |
| Primary tunnel     | LSP BFD                 | 1-2256           | 513-8704        |
| Backup tunnel      | LSP BFD                 | 1-2256           | 8705-16896      |
| MBB primary tunnel | LSP BFD                 | 1-2256           | 16897-25088     |
| MBB backup tunnel  | LSP BFD                 | 1-2256           | 25089-33280     |
| PW                 | VCCV BFD                | Not supported    | 33281-66049     |

## Procedural flow for BFD dampening

This section provides a few flowcharts to help understand how dampening processing works in SAOS.

### Dampening processing for a BFD session during startup and an operating state change from up to down

The following two flowcharts depict this process.



**Figure 17-1**  
Dampening process 1—procedural flow part A





**Figure 17-2**  
**Dampening process 1—procedural flow part B**



### Dampening processing for a BFD session when the operating state changes from down to up and during timer expiry

The following two flowcharts depict this process.

**Figure 17-3**  
Dampening process 2—procedural flow part A




**Figure 17-4**  
**Dampening process 2—procedural flow part B**





**Figure 17-5**  
**Dampening process 2—procedural flow part C**



## Procedures and examples

This chapter provides the following procedures and examples for BFD:

- “Enabling or disabling BFD” on page 17-8
- “Managing BFD profiles” on page 17-10
- “Managing BFD HW acceleration” on page 17-13
- “Setting BFD HW acceleration to default values” on page 17-14
- “Managing VCCV BFD HW acceleration” on page 17-15
- “Setting VCCV BFD HW acceleration to default values” on page 17-16
- “Managing BFD sessions” on page 17-17
- “Displaying BFD session information” on page 17-20
- “Displaying BFD default profiles” on page 17-22
- “Single hop IP-BFD example” on page 17-24
- “Single hop IP-BFD default profile example” on page 17-26
- “Multi-hop IP-BFD default profile example” on page 17-27
- “BFD MPLS-BFD configuration example” on page 17-28

- [“BFD over static MPLS tunnels example” on page 17-29](#)
- [“Dynamic MPLS-TE tunnels example” on page 17-30](#)
- [“Static co-routed MPLS-TP tunnels example” on page 17-31](#)
- [“Static associated MPLS-TP tunnels example” on page 17-32](#)
- [“Dynamic associated MPLS-TP tunnels example” on page 17-33](#)
- [“Static co-routed primary tunnel and dynamic associated backup tunnel example” on page 17-34](#)
- [“Dynamic associated as primary tunnel and static co-routed as backup tunnel example” on page 17-35](#)
- [“Static co-routed primary tunnel and static associated backup tunnel example” on page 17-36](#)
- [“Static associated primary tunnel and static co-routed backup tunnel example” on page 17-37](#)
- [“VCCV BFD over static SS-PW over single LSP example” on page 17-38](#)
- [“VCCV BFD over static MS-PW over single LSP example” on page 17-41](#)
- [“BFD dampening on BFD over static co-routed LSP example” on page 17-45](#)
- [“Displaying BFD dampening profiles” on page 17-48](#)
- [“Displaying the default BFD dampening profile” on page 17-49](#)

## Procedure 17-1

### Enabling or disabling BFD

BFD is globally disabled unless an MPLS license is installed.

You can:

- enable BFD globally
- disable BFD globally
- display the BFD global state

| Step | Action |
|------|--------|
|------|--------|

#### To enable BFD globally

- 1 Enable BFD globally:  
`bfd enable`

#### To disable BFD globally

- 2 Disable BFD globally:  
`bfd disable`

#### To display the BFD global state

- 3 Display the BFD global state.  
`bfd show`

—end—

### Example

The following example shows sample output for the `bfd show` command on a 5142 device.

```
3928*> bfd show
```

```
+-----+
+ BFD GLOBAL INFORMATION +
+-----+
+ Admin State Enabled +
+ LSP HW Acceleration Enabled +
+ VCCV HW Acceleration Enabled +
+ HW Sessions 5 +
+-----+
+ BFD SESSION INFORMATION +
+-----+
+ Sessions Created 5 +
+ Sessions Admin Up 4 +
+ Sessions Oper Up 4 +
+-----+
+ IP BFD SESSION INFORMATION +
+-----+
+ Sessions Created 0 +
+-----+
```

|                                  |     |
|----------------------------------|-----|
| Sessions Admin Up                | 0   |
| Sessions Oper Up                 | 0   |
| +-----+-----+                    |     |
| + LSP BFD SESSION INFORMATION +  |     |
| +-----+-----+                    |     |
| Sessions Created                 | 3   |
| Sessions Admin Up                | 2   |
| Sessions Oper Up                 | 2   |
| +-----+-----+                    |     |
| + VCCV BFD SESSION INFORMATION + |     |
| +-----+-----+                    |     |
| Sessions Created                 | 2   |
| Sessions Admin Up                | 2   |
| Sessions Oper Up                 | 2   |
| +-----+-----+                    |     |
| + BFD PROFILE INFORMATION +      |     |
| +-----+-----+                    |     |
| Profiles Created                 | 12  |
| Profiles in Use                  | 2   |
| Max. Profiles                    | 135 |
| +-----+-----+                    |     |



## Procedure 17-2

### Managing BFD profiles

---

If specified without any parameters, the command creates a profile with default parameters as follows:

- Tx/Rx timer intervals vary as per session type and platform type. LSP BFD has lower timer values compared to VCCV BFD, maintaining a fault hierarchy. Similarly, HW-based sessions have lower time values compared to SW-based sessions. Refer to [“Displaying BFD default profiles” on page 17-22](#) for specific details.

**Note:** The command supports the following fixed configurable BFD Tx/Rx control packet intervals:

- 3.3 msec
- 10 msec
- 20 msec
- 50 msec
- 100 msec
- 300 msec
- 1 sec
- 10 sec
- multiplier 3 (multiplier is the detection time multiplier)
- role active
- G-ACh type
- hw-acceleration

Profile parameters can be modified only if the profile is not in use.

You can

- create a profile
- modify a profile
- return BFD parameters of a specific profile to default values
- delete a profile
- show profiles including the defaults

---

| Step | Action |
|------|--------|
|------|--------|

---

### To create a profile

- 1 Create a profile:  

```
bfd profile create profile <bfd-profile-name> [{transmit-
interval <3.3-10000 ms>} {receive-interval <3.3-
10000msec>}] [role <active|passive>] [lsp-gachtype 7]
[hw-acceleration <disable | enable>]
```

where

profile <bfd-profile-name> is the name of the profile.

{transmit-interval <3.3-10000 msec>} is the transmit control packet interval. The default value is 100 ms.

>

receive-interval <3.3-10000msec> is the receive control packet interval. The default value is 100 ms.

>

role <active|passive> is the role. The default value is active.

lsp-gachtype 7 sets the G-ACh type to 7.

hw-acceleration <disable | enable> enables or disables HW acceleration for BFD. This option is available on 3926m, 3928, 3942, 5142 and 5160 devices.

### To modify a profile

- 2 Modify a profile:  

```
bfd profile set <bfd-profile-name> [{transmit-interval
<3.3-10000msec>} {receive-interval <3.3-10000msec>}]
[role <active|passive>][hw-acceleration <disable |
enable>]
```

### To return BFD parameters of a specific profile to default values

- 3 Return BFD parameters of a specific profile to default values:  

```
bfd profile unset <bfd-profile-name> [transmit-interval |
receive-interval] [role] [lsp-gachtype] [hw-
acceleration]
```

### To delete a profile

- 4 Delete a profile:  

```
bfd profile delete <bfd-profile-name>
```

### To show profiles including the defaults

- 5 Show profiles including the defaults:  

```
bfd profile show [profile <profile_name>]
```

—end—

### Example

The following example creates a BFD profile that uses 300 msec transmit and receive intervals, active role, and async mode.

```
bfd profile create profile bfd_prof1 tx-interval 300 rx-interval 300 role active
```

The following example creates a BFD profile with a G-ACh type of 7 with hw-acceleration enabled.

```
bfd profile create profile bfd_prof1 lsp-gachtype 7 hw-acceleration enable
```

The following example creates a BFD profile that uses 100 ms transmit interval, 300 ms receive interval, passive role, and async mode.

```
bfd profile create profile bfd_profile2 tx-interval 100 rx-interval 300 role passive
```

The following example creates a BFD profile that uses passive role, async mode, and all other parameters are set to default values.

```
bfd profile create profile bfd_profile3 role passive
```

---

## Procedure 17-3

### Managing BFD HW acceleration

---

You can

- set the device in hardware mode
- set the device in software mode

---

| Step | Action |
|------|--------|
|------|--------|

---

|   |                                        |
|---|----------------------------------------|
| 1 | Enable or disable BFD HW acceleration: |
|---|----------------------------------------|

```
bfd set [hw-acceleration <on | off>]
```

where

|                               |                                                                            |
|-------------------------------|----------------------------------------------------------------------------|
| hw-acceleration<br><on   off> | on sets the device in hardware mode. Off sets the device in software mode. |
|-------------------------------|----------------------------------------------------------------------------|

—end—

## Procedure 17-4

### Setting BFD HW acceleration to default values

---

You can return BFD hardware acceleration to default values.

---

| Step | Action |
|------|--------|
|------|--------|

---

- |   |                                                  |
|---|--------------------------------------------------|
| 1 | Set BFD hardware acceleration to default values: |
|---|--------------------------------------------------|

```
bfd unset [hw-acceleration]
```

where

hw-acceleration controls HW acceleration capability for the device.

—end—

## Procedure 17-5

### Managing VCCV BFD HW acceleration

You can set the VCCV device in hardware mode.

| Step | Action |
|------|--------|
|------|--------|

- |   |                                                                                                                                                                                       |
|---|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | <p>Enable VCCV BFD HW acceleration:</p> <pre>bfd set [vccv-hw-acceleration &lt;on&gt;]</pre> <p>where</p> <p>vccv-hw-acceleration &lt;on&gt; on sets the device in hardware mode.</p> |
|---|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

—end—

## Procedure 17-6

### Setting VCCV BFD HW acceleration to default values

---

You can return VCCV BFD HW acceleration to default values.

---

| Step | Action |
|------|--------|
|------|--------|

---

- |   |                                                 |
|---|-------------------------------------------------|
| 1 | Set VCCV BFD HW acceleration to default values: |
|---|-------------------------------------------------|

```
bfd unset [vccv-hw-acceleration]
```

where

|                      |                                                                   |
|----------------------|-------------------------------------------------------------------|
| vccv-hw-acceleration | controls VCCV BFD HW acceleration capability for the VCCV device. |
|----------------------|-------------------------------------------------------------------|

—end—

---

## Procedure 17-7

### Managing BFD sessions

---

You can

- bind the session to the named next hop neighbor or tunnel
- delete a BFD session
- administratively enable a session
- administratively disable a session
- set a session to use values in the (new) specified profile
- unset a session to use values in the (new) specified profile
- display statistics for a specified BFD session
- clear statistics for a specified session

**Note 1:** If no profile is specified, defaults are used.

**Note 2:** If administratively disabled, commands for the session do not take effect and the session does not respond to received BFD packets.

**Note 3:** Changing to a different profile is possible only if the session has not been enabled.

Statistics for a specified session are:

- counts of BFD control packets sent and received
- cumulative number of times the session went down
- current session up time

When you clear statistics, all statistics are cleared.

Fields for a specified session instance are:

- BFD session names
- neighbor IP address
- tunnel name
- administrative and operational states
- myDiscriminator
- yourDiscriminator
- multiplier
- role
- negotiated transmit interval
- negotiated receive interval



- detection time
- profile name
- session Id
- configured transmit and receive intervals

---

| Step | Action |
|------|--------|
|------|--------|

---

***To bind the session to the named next hop neighbor or tunnel***

- 1 Bind the session to the named next hop neighbor or tunnel (static or dynamic):  

```
bfd session create session <session_name> {neighbor <IP address>} [profile <bfd-profile-name>]
```

***To delete a BFD session***

- 2 Delete a BFD session:  

```
bfd session delete session <session_name>
```

where

session                      is the session to be deleted.  
<session\_name>

***To administratively enable a session***

- 3 Administratively enable a session:  

```
bfd session enable session <session_name>
```

where

session                      is the session to be administratively enabled.  
<session\_name>

***To administratively disable a session***

- 4 Administratively disable a session:  

```
bfd session disable session <session_name>
```

where

session                      is the session to be administratively disabled.  
<session\_name>

***To set a session to use values in the (new) specified profile:***

- 5 Set a session to use values in the (new) specified profile:  

```
bfd session set session <session_name> profile
<bfd_profile_name>
```

**To unset a session to use values in the (new) specified profile:**

- 6        Unset a session to use values in the (new) specified profile:  
          bfd session unset session <session\_name> profile  
          where  
          session            is the session to be administratively disabled.  
          <session\_name>

**To display statistics for a specified BFD session**

- 7        Display statistics for a specified BFD session:  
          bfd session show session <session> {statistics}  
          where  
          session            is the BFD session name.  
          <session>  
          statistics        is total statistics.

**To clear statistics for a specified session**

- 8        Clear statistics for a specified session:  
          bfd session clear session <session\_name> statistics  
          where  
          session            is the session that you want to clear statistics for.  
          <session\_name>

—end—

## Example

The following example binds session XYZ to the next-hop neighbor 10.10.10.2.

```
bfd session create session XYZ neighbor 10.10.10.2
bfd session enable session XYZ
```

The following example binds the session ex2 to the next hop neighbor 192.0.2.0 and uses the BFD profile bfd\_prof1.

```
bfd session create session EX2 profile bfd_prof1 neighbor 192.0.2.1
bfd session enable session EX2
```

The following example binds the session IP-BFD3 to the next-hop neighbor 198.51.100.0 and uses bfd profile bfd\_prof100.

```
bfd session create session IP-BFD3 neighbor 198.51.100.1
bfd session set session IP-BFD3 profile bfd_prof100
bfd session enable session IP-BFD3
```

## Procedure 17-8

### Displaying BFD session information

You can display

- a BFD session
- selected information for a BFD session
- PW information for VCCV BFD

| Step | Action |
|------|--------|
|------|--------|

#### To display a BFD session

- 1 Display a BFD session:

```
bfd session show
```

#### To display selected information for a BFD session

- 2 Display selected information for a BFD session:

```
bfd session show {all-ip} {all-lsp} {all-vccv} {neighbor
<IP address>} {oper-down} {oper-up} {statistics} {vc}
{tunnel} {te-assoc <MPLS assoc te-tunnel>} {tp-assoc
<MPLS assoc tp-tunnel>} {tp-egress-corout <MPLS egress
tp-corout-tunnel>} {tp-ingress-corout <MPLS ingress tp-
corouttunnel>}
```

where

|                                 |                                            |
|---------------------------------|--------------------------------------------|
| all-ip                          | specifies all IP BFD sessions.             |
| all-lsp                         | specifies all LSP BFD sessions.            |
| all-vccv                        | specifies all VCCV BFD sessions.           |
| neighbor <IP address>           | is the next-hop IP address.                |
| oper-down                       | specifies all operationally-down sessions. |
| oper-up                         | specifies all operationally-up sessions.   |
| statistics                      | specifies all statistics.                  |
| vc                              | is the PW.                                 |
| tunnel                          | is the tunnel.                             |
| te-assoc <MPLS assoc te-tunnel> | is the list of associated TE tunnels.      |

where

tp-assoc <MPLS is the list of associated TP tunnels.  
assoc tp-tunnel>

tp-egress-corout is the list of TP egress tunnels.  
<MPLS egress  
tp-corout-tunnel>

tp-ingress-corout is the list of TP ingress tunnels.  
<MPLS ingress  
tp-corouttunnel>

—end—

## Example

In this session instance output for VCCV BFD, the PW information is displayed:

```
pns-6-16-5160*> bfd session show session VBFS_0000000205_svc_2>4_4
```

| BFD SESSION         |                           |
|---------------------|---------------------------|
| Session Name        | VBFS_0000000205_svc_2>4_4 |
| Session Index       | 33281                     |
| My Discriminator    | 199262721                 |
| Your Discriminator  | 3500179969                |
| BFD Session Type    | VCCV                      |
| PW INFORMATION      |                           |
| VC Index            | 1001                      |
| PW ID               | 205                       |
| VC Name             | svc_2>4_4                 |
| Incoming Label      | 2005                      |
| Outgoing Label      | 4005                      |
| Operating CC        | PW ACH                    |
| Operating BFD CV    | BFD-Ach-Detect-Only       |
| VC Admin State      | Enabled                   |
| PROFILE INFORMATION |                           |
| Profile Name        | bfdProf2                  |

## Procedure 17-9

### Displaying BFD default profiles

Display the BFD default profiles for each platform.

For platforms supporting hardware acceleration (3926m, 3928, 3942, 5142 and 5160), the BFD default behavior has been updated as follows:

**Table 17-2**  
**BFD defaults for platforms supporting hardware acceleration**

| BFD defaults                           | Before 6.18 | After 6.18              |
|----------------------------------------|-------------|-------------------------|
| Global LSP hardware acceleration       | Disabled    | Enabled                 |
| Global VCCV hardware acceleration      | N/A         | Enabled                 |
| Default LSP BFD profile                | Active-LSP  | Active-LSP-Fast-Default |
| Default VCCV BFD profile               | N/A         | VCCV-Fast-Default       |
| Per profile hardware acceleration      | Disabled    | Enabled                 |
| Per profile transmit/receive intervals | 100 ms      | 100 ms                  |
| Default BFD dampening profile          | N/A         | Dampen-Default          |

For platforms that do not support hardware acceleration, the BFD default behavior is as follows:

**Table 17-3**  
**BFD defaults for platforms not supporting hardware acceleration**

| BFD defaults                           | Before 6.18 | After 6.18        |
|----------------------------------------|-------------|-------------------|
| Global LSP hardware acceleration       | Disabled    | Disabled          |
| Global VCCV hardware acceleration      | N/A         | Disabled          |
| Default LSP BFD profile                | Active-LSP  | Active-LSP        |
| Default VCCV BFD profile               | N/A         | VCCV-Slow-Default |
| Per profile hardware acceleration      | Disabled    | Disabled          |
| Per profile transmit/receive intervals | 100 ms      | 100 ms            |
| Default BFD dampening profile          | N/A         | Dampen-Default    |

| Step | Action |
|------|--------|
|------|--------|

**1** To display BFD default profiles:

```
> bfd profile show
```

—end—

**Example**

```
5160> bfd profile show
```

| BFD PROFILES |                          |                     |                    |         |                 |                       |
|--------------|--------------------------|---------------------|--------------------|---------|-----------------|-----------------------|
| Index        | Profile Name             | Transmit Intvl (ms) | Receive Intvl (ms) | Role    | Sessions in Use | Hardware Acceleration |
| 1            | IP-Default               | 100                 | 100                | Active  | 0               | Off                   |
| 2            | Active-LSP               | 100                 | 100                | Active  | 0               | Off                   |
| 3            | Passive-LSP              | 100                 | 100                | Passive | 0               | Off                   |
| 4            | bfdProf1                 | 3.3                 | 3.3                | Active  | 20              | On                    |
| 129          | IP-IGP-Default           | 100                 | 100                | Active  | 0               | Off                   |
| 130          | IP-BGP-Singlehop-Default | 100                 | 100                | Active  | 0               | Off                   |
| 131          | IP-BGP-Multihop-Default  | 100                 | 100                | Active  | 0               | Off                   |
| 132          | Active-LSP-Fast-Default  | 10                  | 10                 | Active  | 0               | On                    |
| 133          | Passive-LSP-Fast-Default | 10                  | 10                 | Passive | 0               | On                    |
| 134          | VCCV-Slow-Default        | 300                 | 300                | Active  | 0               | Off                   |
| 135          | VCCV-Fast-Default        | 50                  | 50                 | Active  | 0               | On                    |

## Procedure 17-10

### Single hop IP-BFD example

This example shows how a single-hop IP-BFD configuration connects Node A and Node B.



- Node A is configured to create a session with a profile which sets the transmit and receive intervals to 300 ms. The rest of the parameters are defaulted.
- Node B is configured to create a session with a profile which sets transmit interval to 300 ms and receive interval to 300 ms. The rest of the parameters are defaulted.

| Step | Action |
|------|--------|
|------|--------|

#### On Node A

- 1 Create a BFD profile named profA1:  

```
bfd profile create profile profA1 transmit-interval 300
receive-interval 300 role active
```
- 2 Create a BFD session to use profile profA1:  

```
bfd session create session sessA1 neighbor 192.0.2.1
profile profA1
```

#### On Node B

- 3 Create a BFD profile named profB1:  

```
bfd profile create profile profB1 transmit-interval 300
receive-interval 300 role passive
```
- 4 Create a BFD session to use profile profB1:  

```
bfd session create session sessB1 neighbor 198.51.100.1
profile profB1
```

—end—

After the session has come up, you can confirm the running transmit and receive intervals by means of the BFD session show command:

- Node A:

- transmit interval 300ms
  - receive interval 300 ms
- Node B:
  - transmit interval 300 ms
  - receive interval 300 ms



## Procedure 17-11

### Single hop IP-BFD default profile example

This example shows how a single-hop IP-BFD configuration connects Node A and Node B.



- Node A is configured to create a session which uses the default profile.
- Node B is configured to create a session with the sample profile as in [“Single hop IP-BFD example” on page 17-24.](#)

| Step | Action |
|------|--------|
|------|--------|

#### On Node A

- 1 Create a BFD session that uses the default profile:  

```
bfd session create session sessA1 neighbor 192.0.2.1
```

#### On Node B

- 2 Create a BFD session to use profile profB1:  

```
bfd session create session sessB2 neighbor 200.20.20.2
profile profB1
```

—end—

After the session has come up, the running transmit and receive intervals are:

- Node A:
  - transmit interval 300 ms
  - receive interval 300 ms
- Node B:
  - transmit interval 300 ms
  - receive interval 300 ms

## Procedure 17-12

### Multi-hop IP-BFD default profile example


This example shows how to configure a multi-hop IP-BFD session.

| Step | Action                                                                                                                                                                                                 |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Create a multi-hop BFD session that uses the default profile:<br><pre>bfd session create session bfd neighbor 10.10.10.1 source<br/>20.20.20.1 multihop</pre> <p style="text-align: center;">—end—</p> |

## Procedure 17-13

### BFD MPLS-BFD configuration example

This example shows a bidirectional MPLS-BFD configuration.



- On LER A, create a BFD session to use the same profA1 profile from the example above.
- On LER B, create a BFD session with a profile which sets transmit interval to 300 ms and receive interval to 300 ms. The rest of the parameters are defaulted.

| Step | Action |
|------|--------|
|------|--------|

#### On LER A

- 1 Configure MPLS to pair an ingress MPLS tunnel and have BFD monitor it:  

```
mpls tunnel create bidir-ingress-assoc TunnelBidir
forward-tunnel TunnelA reverse-dyntun-name TunnelB bfd-
monitor enable bfd-profile profA1
```

#### On LER B

- 2 Create the BFD profile:  

```
bfd profile create profile pfLERB1 transmit-interval 300
receive-interval 300 role passive
```
- 3 Configure MPLS to pair an ingress to an egress MPLS tunnel and have BFD monitor it such that LER A and LER B form peers of a bidirectional MPLS tunnel:  

```
mpls tunnel create bidir-ingress-assoc TunnelBidir
forward-tunnel TunnelB reverse-dyntun-name TunnelA bfd-
monitor enable bfd-profile pfLERB1
```

—end—

After the session has come up, the running transmit and receive intervals are:

- LER A: transmit interval 300ms, receive interval 300 ms
- LER B: transmit interval 300 ms, receive interval 300 ms

## Procedure 17-14

### BFD over static MPLS tunnels example

This example shows a sample configuration for BFD over static MPLS tunnels.



| Step | Action |
|------|--------|
|------|--------|

#### On Node A

- 1 Configure the ingress:  

```
mpls tunnel create static-ingress static-TE-AtoB dest-ip 198.51.100.1 next-hop-ip 203.0.113.2 out-label 2000
```
- 2 Configure the egress:  

```
mpls tunnel create static-egress static-TE BtoA src-ip 198.51.100.1 in-label 2001
```
- 3 Configure the bidirectional ingress associated:  

```
mpls tunnel create bidir-ingress-assoc static-TE-AtoB-
assoc forwardtunnel static-TE-AtoB reverse-static-tunnel
static-TE-BtoA bfd-monitor enable
```

#### On Node B

- 4 Configure the ingress:  

```
mpls tunnel create static-ingress static-TE-BtoA dest-ip 192.0.2.1 next-hop-ip 203.0.113.1 out-label 2001
```
- 5 Configure the egress:  

```
mpls tunnel create static-egress static-TE-AtoB src-ip 192.0.2.1 in-label 2000
```
- 6 Create the bidirectional ingress associated:  

```
mpls tunnel create birdir-ingress-assoc static-TE-BtoA-
assoc forwardtunnel static-TE-BtoA reverse-static-tunnel
static-TE-AtoB bfd-monitor enable
```

—end—

## Procedure 17-15

### Dynamic MPLS-TE tunnels example

This example shows how to configure dynamic MPLS-TE tunnels.

| Step | Action |
|------|--------|
|------|--------|

**On Node A**

- |   |                                                                                                                                                                                                                          |
|---|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | Configure the rsvp-ingress:<br><pre>mpls tunnel create rsvp-ingress dynamic-TE-AtoB dest-ip 198.51.100.1</pre>                                                                                                           |
| 2 | Configure the bidirectional ingress associated:<br><pre>mpls tunnel create bidir-ingress-assoc dynamic-TE-AtoB-<br/>assoc forwardtunnel dynamic-TE-AtoB reverse-dyntun-name<br/>dynamic-TE-BtoA bfd-monitor enable</pre> |

**On Node B**

- |   |                                                                                                                                                                                                                 |
|---|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3 | Configure the rsvp-ingress:<br><pre>mpls tunnel create rsvp-ingress dynamic-TE-BtoA dest-ip 192.0.2.1</pre>                                                                                                     |
| 4 | Configure bidirectional-ingress-assoc:<br><pre>mpls tunnel create bidir-ingress-assoc dynamic-TE-BtoA-<br/>assoc forwardtunnel dynamic-TE-BtoA reverse-dyntun-name<br/>dynamic-TE-AToB bfd-monitor enable</pre> |

—end—

## Procedure 17-16

### Static co-routed MPLS-TP tunnels example

This example shows how to configure static co-routed MPLS-TP tunnels.

| Step             | Action                                                                                                                                                                                                                          |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>On Node A</b> |                                                                                                                                                                                                                                 |
| 1                | Configure the static ingress co-routed:<br><pre>gmpls tp-tunnel create static-ingress-corout sTP-Corout- AtoB dest-ip 198.51.100.0 next-hop-ip 203.0.113.2 forward-out-label 2200 reverse-inlabel 2201 bfd-monitor enable</pre> |
| <b>On Node B</b> |                                                                                                                                                                                                                                 |
| 2                | Configure the static egress co-routed:<br><pre>gmpls tp-tunnel create static-egress-corout sTP-Corout- AtoB src-ip 192.0.2.1 prev-hop-ip 203.0.113.2 forward-in- label 2200 reverse-out-label 2201 bfd-monitor enable</pre>     |
| —end—            |                                                                                                                                                                                                                                 |

## Procedure 17-17

### Static associated MPLS-TP tunnels example

This example shows how to configure static associated MPLS-TP tunnels.

| Step | Action |
|------|--------|
|------|--------|

#### On Node A

- |   |                                                                                                                                                                                                           |
|---|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | Configure the static ingress unidirectional:<br><pre>gmpls tp-tunnel create static-ingress-unidir sTP-Unidir- AtoB dest-ip 198.51.100.1 next-hop-ip 203.0.113.2 forward-out-label 2210</pre>              |
| 2 | Configure the static egress unidirectional:<br><pre>gmpls tp-tunnel create static-egress-unidir sTP-Unidir- BtoA src-ip 198.51.100.1 prev-hop-ip 203.0.113.1 forward- in-label 2211</pre>                 |
| 3 | Configure the static ingress associated:<br><pre>gmpls tp-tunnel create static-ingress-assoc sTP-BtoA- assoc forwardtunnel sTP-Unidir-AtoB reverse-static- tunnel sTP-Unidir-BtoA bfdmonitor enable</pre> |

#### On Node B

- |   |                                                                                                                                                                                                       |
|---|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 4 | Configure the static ingress unidirectional:<br><pre>gmpls tp-tunnel create static-ingress-unidir sTP-Unidir- BtoA dest-ip 192.0.2.1 next-hop-ip 203.0.113.1 forward- out-label 2211</pre>            |
| 5 | Configure the static egress unidirectional:<br><pre>gmpls tp-tunnel create static-egress-unidir sTP-Unidir- AtoB src-ip 192.0.2.1 forward-in-label 2210</pre>                                         |
| 6 | Configure static-ingress-associated:<br><pre>gmpls tp-tunnel create static-ingress-assoc sTP-AtoB- assoc forwardtunnel sTP-Unidir-BtoA reverse-static- tunnel sTP-Unidir-AtoB bfdmonitor enable</pre> |

—end—

## Procedure 17-18

### Dynamic associated MPLS-TP tunnels example

This example shows how to configure dynamic associated MPLS-TP tunnels.

| Step | Action |
|------|--------|
|------|--------|

#### On Node A

- |   |                                                                                                                                                                                                                   |
|---|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | Configure the rsvp ingress unidirectional:<br><code>gmpls tp-tunnel create rsvp-ingress-unidir dTP-Unidir-AtoB dest-ip 198.51.100.1</code>                                                                        |
| 2 | Configure the static ingress associated:<br><code>gmpls tp-tunnel create static-ingress-assoc dTP-AtoB-<br/>assoc forwardtunnel dTP-Unidir-AtoB reverse-dyntun-name<br/>dTP-Unidir-BtoA bfd-monitor enable</code> |

#### On Node B

- |   |                                                                                                                                                                                                                   |
|---|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3 | Configure the rsvp ingress unidirectional:<br><code>gmpls tp-tunnel create rsvp-ingress-unidir dTP-Unidir-BtoA dest-ip 192.0.2.1</code>                                                                           |
| 4 | Configure the static ingress associated:<br><code>gmpls tp-tunnel create static-ingress-assoc dTP-BtoA-<br/>assoc forwardtunnel dTP-Unidir-BtoA reverse-dyntun-name<br/>dTP-Unidir-AtoB bfd-monitor enable</code> |

—end—



## Procedure 17-19

### Static co-routed primary tunnel and dynamic associated backup tunnel example

This example shows how to configure the static co-routed tunnel as the primary tunnel and the dynamic associated tunnel as the backup tunnel.

| Step | Action                                                                                                                                                                                                                        |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Configure the static ingress co-routed:<br><pre>gmpls tp-tunnel create static-ingress-corout tp3-in-co-1 dest-ip 10.10.10.10 next-hop-ip 166.166.166.100 forward- out-label 251 reverse-in-label 252 bfd-monitor enable</pre> |
| 2    | Configure the rsvp ingress unidirectional:<br><pre>gmpls tp-tunnel create rsvp-ingress-unidir tp3-rsvp-1 dest-ip 10.10.10.10 backup-tunnel tp3-in-co-1</pre>                                                                  |
| 3    | Configure the static ingress associated:<br><pre>gmpls tp-tunnel create static-ingress-assoc tp-assoc-1 forward-tunnel tp3-rsvp-1 reverse-dyntun-name tp1-rsvp-1 bfd-monitor enable</pre>                                     |

—end—

## Procedure 17-20

### Dynamic associated as primary tunnel and static co-routed as backup tunnel example

This example shows how to configure the dynamic associated tunnel as the primary tunnel and the static-co-routed as the backup tunnel.

| Step | Action                                                                                                                                                                                                                                                                                          |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Configure the rsvp ingress unidirectional:<br><pre>gmpls tp-tunnel create rsvp-ingress-unidir tp1-rsvp-1 dest-ip 40.40.40.40</pre>                                                                                                                                                              |
| 2    | Configure the static ingress associated:<br><pre>gmpls tp-tunnel create static-ingress-assoc tp-assoc-1 forward-tunnel tp-rsvp-1 reverse-dyntun-name tp-rsvp-3 bfd-monitor enable</pre>                                                                                                         |
| 3    | Configure the static ingress co-routed:<br><pre>gmpls tp-tunnel create static-egress-corout tp-eg-co-1- bak src-ip 40.40.40.40 prev-hop-ip 166.166.166.200 forward-in-label 271 reverse-out-label 272 backup-tunnel tp-rsvp-1 bfd-monitor enable</pre> <p style="text-align: center;">—end—</p> |

## Procedure 17-21

### Static co-routed primary tunnel and static associated backup tunnel example

This example shows how to configure the static co-routed tunnel as the primary tunnel and the static associated tunnel as the backup tunnel.

| Step | Action                                                                                                                                                                                                                                                                                                                                        |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Configure the static ingress co-routed tunnel:<br><pre>gmpls tp-tunnel create static-ingress-corout tp3-in-co-1 dest-ip 10.10.10.10 next-hop-ip 166.166.166.100 forward- out-label 251 reverse-in-label 252 bfd-monitor enable</pre>                                                                                                          |
| 2    | Configure the static unidirectional tunnels:<br><pre>gmpls tp-tunnel create static-ingress-unidir tp3-st-1 dest-ip 10.10.10.10 next-hop-ip 10.10.1.2 forward-out- label 261 backup tunnel tp3-in-co-1  gmpls tp-tunnel create static-egress-unidir tp-1-st-1 logical-id 1 src-ip 10.10.10.10 forward-in-label 262 prev-hop-ip 10.10.1.2</pre> |
| 3    | Configure the static ingress associated tunnel:<br><pre>gmpls tp-tunnel create static-ingress-assoc tp-assoc- forward-tunnel tp3-st-1 reverse static-tunnel tp-1-st-1 bfd-monitor enable</pre> <p style="text-align: center;">—end—</p>                                                                                                       |

## Procedure 17-22

### Static associated primary tunnel and static co-routed backup tunnel example

This example shows how to configure the static associated tunnel as the primary tunnel and the static co-routed tunnel as the backup tunnel.


| Step | Action                                                                                                                                                                                                                                                                                                                |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Configure the static unidirectional tunnels:<br><pre> gmppls tp-tunnel create static-ingress-unidir tp1-st-3 next-hop-ip 166.166.166.100 dest-ip 10.10.10.10 forward- out-label 271  gmppls tp-tunnel create static-egress-unidir tp-st-3 src- ip 10.10.10.10 forward-in-label 272 prev-hop-ip 166.166.166.100 </pre> |
| 2    | Configure the static-egress co-routed tunnels:<br><pre> gmppls tp-tunnel create static-egress-corout tp1-st-3 bak dest-ip 10.10.10.10 next-hop-ip 10.10.1.2 forward-out- label 281 reverse-in-label 282 bfd-monitor enable backup- tunnel tp1-st-3 </pre>                                                             |
| 3    | Configure the static ingress associated tunnel:<br><pre> gmppls tp-tunnel create static-ingress-assoc tp-assoc-3 forward-tunnel tp1-st-3 reverse-static-tunnel tp-st-3 bfd-monitor enable </pre> <p style="text-align: center;">—end—</p>                                                                             |

## Procedure 17-23

### VCCV BFD over static SS-PW over single LSP example

This figure illustrates a topology for VCCV BFD over static SS-PW over single LSP:

**Figure 17-6**  
VCCV BFD over static SS-PW over single LSP topology



The co-routed tunnel is configured over IPv4 interfaces and non-LAG ports. The tunnel must be a static co-routed tunnel (as in the example).

| Step | Action |
|------|--------|
|------|--------|

#### On PE-1

- Configure the static ingress co-routed tunnel via LSR-1:  

```
gmpls tp-tunnel create static-ingress-corout scrt_1->3_2
dest-ip 3.3.3.3 next-hop-ip 111.0.0.2 forward-out-label
101 reverse-in-label 201
```
- Create the BFD profile:  

```
bfd profile create profile bfdProf1 receive-interval
50msec transmit-interval 50msec role active
hw-acceleration enable
```

**Note:** HW acceleration is available only on platforms where BFD support is natively available on the hardware. For details, see the table titled [Maximum VCCV BFD sessions](#).

- 3 Create the VCCV profile with the CV type and CC type:  

```
mpls l2-vpn pw-vccv-profile create vccv-profile
vccvProfCw cc-in-band on cc-ttl-exp off cc-ciena-oob off
cv-bfd-ach-detect-only on
```

**Note:** For VCCV-BFD, CC-Type 1 (cc-in-band) and CC-Type 4 (cc-ciena-oob) are supported. CC-Type 1 must be used if the PW is an Infrastructure PW. CC-Type 4 is supported only for BFD sessions in software.
- 4 Configure global BFD:  

```
bfd set hw-acceleration on
bfd enable
```
- 5 Configure static PW (with CW support):  

```
mpls l2-vpn create static-vc svc_1->3_2 pw-id 201 peer
3.3.3.3 in-label 2001 out-label 1001 tp-tunnel-ingr-
corout scrt_1->3_2 pw-cword on pw-vccv-profile vccvProfCw
status-tlv on
```
- 6 Configure the VS:  

```
virtual-switch ethernet create vs vsSt2 mode vpws
virtual-switch ethernet attach mpls-vc svc_1->3_2 vs
vsSt2
```
- 7 Configure and enable BFD on the PW:  

```
mpls l2-vpn set static-vc svc_1->3_2 bfd-monitor enable
bfd-profile bfdProf1
```

## On PE-2

- 1 Configure the static egress co-routed tunnel via LSR-1:  

```
gmpls tp-tunnel create static-egress-corout scrt_1->3_2
src-ip 1.1.1.1 prev-hop-ip 121.0.0.1 forward-in-label 301
reverse-out-label 401
```
- 2 Create the BFD profile:  

```
bfd profile create profile bfdProf1 receive-interval
50msec transmit-interval 50msec role active
hw-acceleration enable
```

**Note:** HW acceleration is available only on platforms where BFD support is natively available on the hardware. For details, see the table titled [Maximum VCCV BFD sessions](#).
- 3 Create the VCCV profile with the CV type and CC type:  

```
mpls l2-vpn pw-vccv-profile create vccv-profile
vccvProfCw cc-in-band on cc-ttl-exp off cc-ciena-oob off
cv-bfd-ach-detect-only on
```

**Note:** For VCCV-BFD, CC-Type 1 (cc-in-band) and CC-Type 4 (cc-ciena-oob) are supported. CC-Type 1 must be used if the PW is an Infrastructure PW. CC-Type 4 is supported only for BFD sessions in software.

- 4**      **Configure global BFD:**  
bfd set hw-acceleration on  
bfd enable
- 5**      **Configure static PW (with CW support):**  
mpls l2-vpn create static-vc svc\_1->3\_2 pw-id 202 peer  
1.1.1.1 tp-tunnel-egrs-corout-static scrt\_1->3\_2 in-  
label 1001 out-label 2001 pw-cword on pw-ccv-profile  
ccvProfCw status-tlv on
- 6**      **Configure the VS:**  
virtual-switch ethernet create vs vsSt2 mode vpws  
virtual-switch ethernet attach mpls-vc svc\_1->3\_2 vs  
vsSt2
- 7**      **Configure and enable BFD on the PW:**  
mpls l2-vpn set static-vc svc\_1->3\_2 bfd-monitor enable  
bfd-profile bfdProf2


—end—

## Procedure 17-24

### VCCV BFD over static MS-PW over single LSP example

This figure illustrates a topology for VCCV BFD over static MS-PW over single LSP:

**Figure 17-7**  
VCCV BFD over static MS-PW over single LSP topology



The co-routed tunnel is configured over IPv4 interfaces and non-LAG ports. The tunnel must be a static co-routed tunnel (as in the example).

| Step | Action |
|------|--------|
|------|--------|

#### On TPE-1

- Configure the static ingress co-routed tunnel via SPE-1:  

```
gmpls tp-tunnel create static-ingress-corout scrt_1->2_1A
dest-ip 2.2.2.2 next-hop-ip 111.0.0.2 forward-out-label
100 reverse-in-label 200
```
- Create the BFD profile:  

```
bfd profile create profile bfdProf1 receive-interval
50msec transmit-interval 50msec role active
hw-acceleration enable
```

**Note:** HW acceleration is available only on platforms where BFD support is natively available on the hardware. For details, see the table titled [Maximum VCCV BFD sessions](#).



- 3 Create the VCCV profile with the CV type and CC type:  

```
mpls l2-vpn pw-vccv-profile create vccv-profile
vccvProfCw cc-in-band on cc-ttl-exp off cc-ciena-oob off
cv-bfd-ach-detect-only on
```

**Note:** For VCCV-BFD, CC-Type 1 (cc-in-band) and CC-Type 4 (cc-ciena-oob) are supported. CC-Type 1 must be used if the PW is an Infrastructure PW. CC-Type 4 is supported only for BFD sessions in software.
- 4 Configure global BFD:  

```
bfd set hw-acceleration on
bfd enable
```
- 5 Configure static PW (with CW support):  

```
mpls l2-vpn create static-vc svc_1->2_1A pw-id 101 peer
2.2.2.2 in-label 2000 out-label 1000 tp-tunnel-ingr-
corout scrt_1->2_1A status-tlv on pw-cword on pw-vccv-
profile vccvProfCw
```
- 6 Configure the VS:  

```
virtual-switch ethernet create vs vsSt1 mode vpws
virtual-switch ethernet attach mpls-vc svc_1->2_1A vs
vsSt1
```
- 7 Configure and enable BFD on the PW:  

```
mpls l2-vpn set static-vc svc_1->2_1A bfd-monitor enable
bfd-profile bfdProf1
```

## On SPE-1

- 1 Configure the static egress co-routed tunnel towards TPE-1 and the static ingress co-routed tunnel towards TPE-2:  

```
gmpls tp-tunnel create static-egress-corout scrt_1->2_1A
src-ip 1.1.1.1 prev-hop-ip 111.0.0.1 forward-in-label 100
reverse-out-label 200
```

```
gmpls tp-tunnel create static-ingress-corout scrt_2->3_1B
dest-ip 3.3.3.3 next-hop-ip 121.0.0.2 forward-out-label
300 reverse-in-label 400
```
- 2 Create the VCCV profile with the CV type and CC type (cc-in-band on, which means CC-Type 1):  

```
mpls l2-vpn pw-vccv-profile create vccv-profile
vccvProfCw cc-in-band on cc-ttl-exp off cc-ciena-oob off
cv-bfd-ach-detect-only on
```

### 3 Configure static PW (with CW support):

```
mpls l2-vpn create static-vc svc_1->2_1A pw-id 101 peer
1.1.1.1 in-label 1000 out-label 2000 tp-tunnel-egrs-
corout-static scrt_1->2_1A pw-mode switching status-tlv
on pw-cword on pw-vccv-profile vccvProfCw

mpls l2-vpn create static-vc svc_2->3_1B pw-id 102 peer
3.3.3.3 in-label 4000 out-label 3000 tp-tunnel-ingr-
corout scrt_2->3_1B pw-mode switching status-tlv on pw-
cword on pw-vccv-profile vccvProfCw
```

### 4 Configure the stitching VS:

```
virtual-switch ethernet create vs mspw-spe-1 mode vpws
virtual-switch ethernet attach mpls-vc svc_1->2_1A vs
mspw-spe-1

virtual-switch ethernet attach mpls-vc svc_2->3_1B vs
mspw-spe-1
```

## On TPE-2

### 1 Configure the static egress co-routed tunnel via SPE-1:

```
gmpls tp-tunnel create static-egress-corout scrt_2->3_1B
src-ip 2.2.2.2 prev-hop-ip 121.0.0.1 forward-in-label 300
reverse-out-label 400
```

### 2 Create the BFD profile:

```
bfd profile create profile bfdProf1 receive-interval
50msec transmit-interval 50msec role active
hw-acceleration enable
```

**Note:** HW acceleration is available only on platforms where BFD support is natively available on the hardware. For details, see the table titled [Maximum VCCV BFD sessions](#).

### 3 Create the VCCV profile with the CV type and CC type:

```
mpls l2-vpn pw-vccv-profile create vccv-profile
vccvProfCw cc-in-band on cc-ttl-exp off cc-ciena-oob off
cv-bfd-ach-detect-only on
```

**Note:** For VCCV-BFD, CC-Type 1 (cc-in-band) and CC-Type 4 (cc-ciena-oob) are supported. CC-Type 1 must be used if the PW is an Infrastructure PW. CC-Type 4 is supported only for BFD sessions in software.

### 4 Configure global BFD:

```
bfd set hw-acceleration on
bfd enable
```

### 5 Configure static PW (with CW support):

```
mpls l2-vpn create static-vc svc_2->3_1B pw-id 102 peer
2.2.2.2 tp-tunnel-egrs-corout-static scrt_2->3_1B in-
label 3000 out-label 4000 status-tlv on pw-cword on pw-
vccv-profile vccvProfCw
```

**6** Configure the VS:

```
virtual-switch ethernet create vs vsSt1 mode vpws
virtual-switch ethernet attach mpls-vc svc_2->3_1B vs
vsSt1
```

**7** Configure and enable BFD on the PW:

```
mpls l2-vpn set static-vc svc_2->3_1B bfd-monitor enable
bfd-profile bfdProf1
```


—end—

## Procedure 17-25

### BFD dampening on BFD over static co-routed LSP example

This figure illustrates a topology for BFD dampening on BFD over static co-routed LSP:

**Figure 17-8**  
BFD dampening on BFD over static co-routed LSP topology



| Step | Action |
|------|--------|
|------|--------|

#### On LER-1

- 1 Configure the unprotected static co-routed ingress tunnel through LSR-1:  

```
gmpls tp-tunnel create static-ingress-corout scrt_1->3_1P
dest-ip 3.3.3.3 next-hop-ip 111.0.0.2 forward-out-label
101 reverse-in-label 201
```
- 2 Create a custom BFD profile:  

```
bfd profile create profile bfdProf3.3 receive-interval
3.3msec transmit-interval 3.3msec hw-acceleration enable
```
- 3 Enable BFD hardware acceleration globally:  

```
bfd set hw-acceleration on
```
- 4 Enable the BFD feature globally:  

```
bfd enable
```

- 5** Associate the custom BFD profile and enable BFD monitoring on the tunnel:  
`gmpls tp-tunnel set static-ingress-corout scrt_1->3_1P  
bfd-profile bfdProf3.3 bfd-monitor enable`
- 6** Create a custom BFD dampening profile:  
`bfd dampening profile create profile lspBfdDampProf_1  
suppress-threshold 250 reuse-threshold 20 decay-half-life  
2000 max-suppress-time 30000 mode full`
- 7** Enable the BFD dampening feature globally:  
`bfd dampening enable`
- 8** Associate the BFD dampening profile and enable dampening on the BFD session corresponding to the tunnel:  
`gmpls tp-tunnel set static-ingress-corout scrt_1->3_1P  
bfd-dampening-profile lspBfdDampProf_1 bfd-dampening on`
- 9** Disable dampening on the LSP BFD session:  
`gmpls tp-tunnel set static-ingress-corout scrt_1->3_1P  
bfd-dampening off`
- 10** Associate a new dampening-profile & re-enable dampening on BFD session corresponding to the tunnel:  
`gmpls tp-tunnel set static-ingress-corout scrt_1->3_1P  
bfd-dampening-profile lspBfdDampProf_2 bfd-dampening on`

**On LSR-1**

- 1** Configure an unprotected static co-routed transit tunnel with ingress as LER-1 and egress as LER-2:  
`gmpls tp-tunnel create static-transit-corout scrt_1->3_1P  
src-ip 1.1.1.1 prev-hop-ip 111.0.0.1 forward-in-label 101  
reverse-out-label 201 dest-ip 3.3.3.3 next-hop-ip  
121.0.0.2 forward-out-label 301 reverse-in-label 401`

**On LER-2**

- 1** Configure the unprotected static co-routed egress tunnel through LSR-1:  
`gmpls tp-tunnel create static-egress-corout scrt_1->3_1P  
src-ip 1.1.1.1 prev-hop-ip 121.0.0.1 forward-in-label 301  
reverse-out-label 401`
- 2** Create a custom BFD profile:  
`bfd profile create profile bfdProf3.3 receive-interval  
3.3msec transmit-interval 3.3msec hw-acceleration enable`
- 3** Enable BFD hardware acceleration globally:  
`bfd set hw-acceleration on`
- 4** Enable the BFD feature globally:  
`bfd enable`

- 5 Associate the custom BFD profile and enable BFD monitoring on the tunnel:  

```
gmpls tp-tunnel set static-egress-corout scrt_1->3_1P
bfd-profile bfdProf3.3 bfd-monitor enable
```
- 6 Create a BFD dampening profile identical to LER-1:  

```
bfd dampening profile create profile lspBfdDampProf_1
suppress-threshold 250 reuse-threshold 20 decay-half-life
2000 max-suppress-time 30000 mode full
```
- 7 Enable the BFD dampening feature globally:  

```
bfd dampening enable
```
- 8 Associate the BFD dampening profile and enable dampening on the BFD session corresponding to the tunnel:  

```
gmpls tp-tunnel set static-egress-corout scrt_1->3_1P
bfd-dampening-profile lspBfdDampProf_1 bfd-dampening on
```
- 9 Disable dampening on the LSP BFD session:  

```
gmpls tp-tunnel set static-ingress-corout scrt_1->3_1P
bfd-dampening off
```
- 10 Associate a new dampening-profile & re-enable dampening on BFD session corresponding to the tunnel:  

```
gmpls tp-tunnel set static-egress-corout scrt_1->3_1P
bfd-dampening-profile lspBfdDampProf_2 bfd-dampening on
```

—end—

## Procedure 17-26

### Displaying BFD dampening profiles

Use this procedure to display BFD dampening profiles.

| Step | Action                                                                                               |
|------|------------------------------------------------------------------------------------------------------|
| 1    | To display a specific BFD dampening profiles:<br>> bfd dampening profile show profile <profile-name> |
| 2    | To display all BFD dampening profiles:<br>> bfd dampening profile show<br>—end—                      |

### Example

In this example, no custom profiles have been configured, so only the default BFD dampening profile is shown.

```
3928*> bfd dampening profile show
```

| BFD DAMPENING PROFILES |                |      |                       |                    |                              |                                |                    |
|------------------------|----------------|------|-----------------------|--------------------|------------------------------|--------------------------------|--------------------|
| Index                  | Profile Name   | Mode | Suppress<br>Threshold | Reuse<br>Threshold | Decay<br>Half Life<br>(msec) | Max<br>Suppress<br>Time (msec) | Profiles<br>in Use |
| 33                     | Dampen-Default | test | 400                   | 15                 | 1000                         | 5000                           | 1                  |

## Procedure 17-27

### Displaying the default BFD dampening profile

Use this procedure to display the default BFD dampening profile.

| Step | Action |
|------|--------|
|------|--------|

- |   |                                                                                                                                                             |
|---|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | To display the default BFD dampening profile:<br><pre>&gt; bfd dampening profile show profile Dampen-Default</pre> <p style="text-align: center;">—end—</p> |
|---|-------------------------------------------------------------------------------------------------------------------------------------------------------------|

### Example

```
3928*> bfd dampening profile show profile Dampen-Default
```

```
3928*> bfd dampening profile show profile Dampen-Default
```

```
+----- BFD DAMPENING PROFILE INFORMATION -----+
Profile Name	Dampen-Default
Suppress-Threshold	400
Reuse-Threshold	15
Max-Suppress-Time(ms)	5000
Decay-Half-Life(ms)	1000
Mode	test
Max Penalty Calculated	480
+-----+-----+
```



## Procedure 17-28

### Managing BFD dampening profiles

---

Use this procedure to manage BFD dampening profiles.

You can

- create a BFD dampening profile
- modify a BFD dampening profile

**Note:** You can modify a BFD dampening profile only if it is not associated with any BFD session.

- return the parameters of a specific BFD dampening profile to default values
- delete a BFD dampening profile
- show a BFD dampening profile, including default parameters

---

| Step | Action |
|------|--------|
|------|--------|

---

### To create a BFD dampening profile

#### 1 Create a BFD dampening profile:

```
bfd dampening profile create profile <bfd-dampening-
profile-name> [{suppress-threshold <(NUMBER:
1..120000)>}] {reuse-threshold <(NUMBER: 1..120000)}]
[decay-half-life <MILLISECONDS: 500..1,800,000>]
[max-suppress-time <MILLISECONDS: 501..1,800,000>]
[mode <test | full>]
```

#### where

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| suppress-threshold | is a value in msec. It specifies the threshold value moving above which notifications/traps are suppressed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| reuse-threshold    | is a value in msec. It specifies the threshold value moving below which notifications/traps are re-enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| decay-half-life    | is a value in msec. It specifies the waiting time after which “penalty” is halved, provided session state remains stable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| max-suppress-time  | is a value in msec. It specifies the maximum time period for which notifications/traps can be suppressed. After this period, notifications/traps are re-enabled irrespective of the session stability or state.                                                                                                                                                                                                                                                                                                                                                                                        |
| mode               | <p>“Test” mode is intended to be used only for getting the results of BFD dampening behavior in a topology, as per user configured dampening parameters, without impacting normal BFD notifications. “Test” mode calculates and stores the results for a given set of configured dampening parameters over a time period, without doing any suppression of the sessions. In “Test” mode, trap and syslog “BFD State Flap Dampening” are not generated.</p> <p>“Full” mode supports suppressing and re-enabling sessions. In “Test” mode, trap and syslog “BFD State Flap Dampening” are generated.</p> |

### To modify a BFD dampening profile

#### 2 Modify a BFD dampening profile:

```
bfd dampening profile set <bfd-dampening-profile-name>
[{suppress-threshold <(NUMBER: 1..120000)>}] {reuse-
threshold <(NUMBER: 1..120000)}] [decay-half-life
<MILLISECONDS: 500..1,800,000>]
[max-suppress-time <MILLISECONDS: 501..1,800,000>]
[mode <test | full>]
```

where

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| suppress-threshold | is a value in msec. It specifies the threshold value moving above which notifications/traps are suppressed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| reuse-threshold    | is a value in msec. It specifies the threshold value moving below which notifications/traps are re-enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| decay-half-life    | is a value in msec. It specifies the waiting time after which “penalty” is halved, provided session state remains stable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| max-suppress-time  | is a value in msec. It specifies the maximum time period for which notifications/traps can be suppressed. After this period, notifications/traps are re-enabled irrespective of the session stability or state.                                                                                                                                                                                                                                                                                                                                                                                 |
| mode               | “Test” mode is intended to be used only for getting the results of BFD dampening behavior in a topology, as per user configured dampening parameters, without impacting normal BFD notifications. “Test” mode calculates and stores the results for a given set of configured dampening parameters over a time period, without doing any suppression of the sessions. In “Test” mode, trap and syslog “BFD State Flap Dampening” are not generated.<br><br>“Full” mode supports suppressing and re-enabling sessions. In “Test” mode, trap and syslog “BFD State Flap Dampening” are generated. |

***To return parameters of a specific BFD dampening profile to default values***

- 3      Return parameters of a BFD dampening profile to the default values:  

```
bfd dampening profile unset profile
<bfd-dampening-profile-name> [suppress-threshold]
[reuse-threshold] [decay-half-life] [max-suppress-time]
[mode]
```

***To delete a BFD dampening profile***

- 4      Delete a BFD dampening profile:  

```
bfd dampening profile delete profile
<bfd-dampening-profile-name>
```

***To show a BFD dampening profile including the defaults***

- 5      Show a BFD dampening profile including the defaults:  

```
bfd dampening profile show profile
<bfd-dampening-profile-name>

—end—
```

---

# Alarm Indication Signal with Link Down Indication

---

This chapter provides the following procedures and examples for Alarm Indication Signal with Link Down Indication (AIS/LDI):

- “Displaying AIS/LDI global state” on page 18-2
- “Enabling and disabling AIS/LDI global state” on page 18-3
- “Displaying AIS/LDI profiles” on page 18-4
- “Configuring AIS/LDI profiles” on page 18-5
- “Configuring AIS/LDI on a static ingress associated TP tunnel (on an LER)” on page 18-9
- “Configuring AIS/LDI on a static transit unidirectional TP tunnel (on an LSR)” on page 18-11
- “Configuring AIS/LDI on a static ingress co-routed TP tunnel (on an LER)” on page 18-12
- “Configuring AIS/LDI on a static egress co-routed TP tunnel (on an LER)” on page 18-13
- “Configuring AIS/LDI on a static transit co-routed TP tunnel (on an LSR)” on page 18-14
- “Displaying AIS/LDI global statistics” on page 18-16
- “Clearing AIS/LDI global statistics” on page 18-18

This chapter provides the following procedures and examples for AIS/LDI:

- “Example topology: AIS over static MPLS-TP associated tunnels” on page 18-19
- “Example topology: AIS over static MPLS-TP co-routed tunnels” on page 18-22

## Procedure 18-1

### Displaying AIS/LDI global state

---

You can display the current global state of AIS/LDI.

---

| Step | Action |
|------|--------|
|------|--------|

---

|   |                                      |
|---|--------------------------------------|
| 1 | Display the global state of AIS/LDI: |
|---|--------------------------------------|

```
ais show
```

—end—

### Example

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

```
> ais show
```

```
+-----+
+ AIS GLOBAL PARAMETERS +
+-----+
| Admin State | Disabled |
+-----+
```

## Procedure 18-2

### Enabling and disabling AIS/LDI global state

You can:

- globally disable AIS/LDI messaging. When AIS/LDI is globally disabled, the system does not send or receive AIS/LDI messages.
- globally enable AIS/LDI messaging. When AIS/LDI is globally enabled, AIS/LDI must also be configured on an MPLS tunnel before the system will send or receive AIS/LDI messages.

By default, AIS/LDI is globally disabled.

| Step | Action |
|------|--------|
|------|--------|

#### To enable AIS/LDI globally

- 1 Enable AIS/LDI globally:  

```
ais enable
```

#### To disable AIS/LDI globally

- 2 Disable AIS/LDI globally:  

```
ais disable
```

—end—

### Example

```
> ais show
+-----+
+ AIS GLOBAL PARAMETERS +
+-----+
+ Admin State | Disabled |
+-----+
> ais enable
> ais show
+-----+
+ AIS GLOBAL PARAMETERS +
+-----+
+ Admin State | Enabled |
+-----+
> ais disable
> ais show
+-----+
+ AIS GLOBAL PARAMETERS +
+-----+
+ Admin State | Disabled |
+-----+
```

## Procedure 18-3

### Displaying AIS/LDI profiles

You can display information about existing AIS/LDI profiles. For each profile, the output displays the:

- index
- profile name
- refresh timer interval in seconds
- AIS RDI state of the profile
- number of LSPs that are using the profile
- mode
- status of signal degrade support

| Step | Action |
|------|--------|
|------|--------|

- |   |                                                     |
|---|-----------------------------------------------------|
| 1 | Display all AIS/LDI profiles or a specific profile. |
|---|-----------------------------------------------------|

```
ais profile show [profile <profile-name>]
```

where

<profile-name> Specifies the name of the profile for to display.

—end—

### Example

The following example shows sample output for the `ais profile show` command.

```
> ais profile show
```

| --AIS PROFILES-- |                |                  |               |            |             |                         |
|------------------|----------------|------------------|---------------|------------|-------------|-------------------------|
| Index            | Profile Name   | Refresh Timer(s) | AIS RDI State | No: Of Lsp | Mode        | Service Degrade Support |
| 1                | Default        | 10               | Enabled       | 0          | Single Byte | Disabled                |
| 2                | SD-AIS-profile | 10               | Enabled       | 0          | Single Byte | Enabled                 |

```
> ais profile show profile SD-AIS-profile
```

| --AIS PROFILES-- |                |                  |               |            |             |                         |
|------------------|----------------|------------------|---------------|------------|-------------|-------------------------|
| Index            | Profile Name   | Refresh Timer(s) | AIS RDI State | No: Of Lsp | Mode        | Service Degrade Support |
| 2                | SD-AIS-profile | 10               | Enabled       | 0          | Single Byte | Enabled                 |

## Procedure 18-4

### Configuring AIS/LDI profiles

You can create, modify or delete AIS/LDI profiles.

**Note 1:** An AIS/LDI profile in use by a tunnel cannot be modified. It can be modified only after it is detached from the tunnel.

**Note 2:** If service degrade is enabled for the MPLS tunnels, signal degrade detection should be enabled on the underlying physical port. See *39XX/51XX Advanced OAM Configuration* for more information.

| Step | Action |
|------|--------|
|------|--------|

#### To create an AIS/LDI profile

- 1 Create an AIS/LDI profile:

```
ais profile create profile <profile-name>
{[ais-rdi <disable|enable>]
[refresh-timer <SECONDS: 1..20>][tlv-mode <single-
byte|double-byte>] [service-degrade] <enable|disable>}
```

where

|                                        |                                                                                                                                                                                             |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <profile-name>                         | Specifies the name of the profile to create.                                                                                                                                                |
| [ais-rdi <disable enable>]             | Specifies whether the AIS-RDI message will be sent by an LER configured with this profile. Default is enabled.                                                                              |
| [refresh-timer <SECONDS: 1..20>]       | Specifies the interval at which AIS/LDI messages are sent.                                                                                                                                  |
| [tlv-mode <single-byte   double-byte>] | Specifies the ais tlv size mode.                                                                                                                                                            |
| [service-degrade] <enable   disable>   | Optional feature that enables or disables service degrade for MPLS services. The default is disabled. If no settings are specified, service-degrade is disabled on the created AIS profile. |

For example:

```
ais profile create profile ais2 refresh-timer 10 tlv-mode single-byte service-
degrade enable
```

```
ais profile show profile ais1
```

| --AIS PROFILES-- |              |                   |               |             |             |                         |
|------------------|--------------|-------------------|---------------|-------------|-------------|-------------------------|
| Index            | Profile Name | Refresh Timer (s) | AIS RDI State | No. Of Lsps | Mode        | Service Degrade Support |
| 3                | ais2         | 10                | Enabled       | 1           | Single Byte | Enabled                 |



## 18-6 Alarm Indication Signal with Link Down Indication

If no parameters are specified, the profile is created with the following defaults:

- ais-rdi: Enabled.
- refresh-timer: 10 seconds.
- Single Byte TLV
- service-degrade disable

For example:

```
ais profile create profile ais2
```

```
ais profile show profile ais2
```

| --AIS PROFILES-- |              |                     |                  |             |             |                            |
|------------------|--------------|---------------------|------------------|-------------|-------------|----------------------------|
| Index            | Profile Name | Refresh<br>Timer(s) | AIS RDI<br>State | No: Of Lsps | Mode        | Service Degrade<br>Support |
| 3                | ais2         | 10                  | Enabled          | 1           | Single Byte | Enabled                    |

### To modify parameters of an existing AIS/LDI profile

#### 2 Modify parameters of an AIS/LDI profile:

```
ais profile set profile <profile-name> {[ais-
rdi<disable|enable>] [refresh-timer<SECONDS:
1..20>][tlv-mode <single-byte | double-byte>][signal-
degrade] <enable | disable>}
```

where

|                                        |                                                                                                                                                                  |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <profile-name>                         | Specifies the name of the profile to modify.                                                                                                                     |
| [ais-rdi<disable enable>]              | Specifies whether the AIS-RDI message will be sent by an LER configured with this profile. Default is enabled.                                                   |
| [refresh-timer<SECONDS: 1..20>]        | Specifies the interval at which AIS/LDI messages are sent.<br><br>A refresh timer configuration can be changed only after the profile is detached from a tunnel. |
| [tlv-mode <single-byte   double-byte>] | Specifies the AIS TLV size mode.                                                                                                                                 |
| [service-degrade] <enable   disable>s  | Optional feature that enables or disables service degrade for MPLS services.                                                                                     |

For example:

```
ais profile set ais1 refresh-timer 10 tlv-mode single-byte enable service-
degrade disable
```

Verify that the parameters were modified:

```
ais profile show
```

| --AIS PROFILES-- |              |                  |               |             |             |                         |
|------------------|--------------|------------------|---------------|-------------|-------------|-------------------------|
| Index            | Profile Name | Refresh Timer(s) | AIS RDI State | No: Of Lsps | Mode        | Service Degrade Support |
| 1                | Default      | 10               | Enabled       | 1           | Single Byte | Disabled                |
| 2                | ais1         | 10               | Enabled       | 0           | Single Byte | Disabled                |

### 3 Set the specified parameters of an AIS/LDI profile to their default values:

```
ais profile unset profile <profile-name>
{[ais-rdi] [refresh-timer] [tlv-mode] [service-degrade]
<enable|disable>}
```

where

<profile-name> Specifies the name of the profile for which values will be returned to default.

ais-rdi Sets ais-rdi to the default value: Enabled.

refresh-timer Sets the refresh-timer to its default value: 10 seconds.

A refresh timer configuration can be changed only after the profile is detached from a tunnel.

tlv-mode Specifies the AIS TLV size mode.

[service-degrade] Optional feature that enables or disables service degrade for MPLS services.

<enable | disable>s

For example, unset the service degrade parameter back to its default of disabled:

```
ais profile unset profile ais1 refresh-timer tlv-mode service-degrade
```

Verify that the AIS/LDI profile now uses default values:

```
ais profile show
```

| --AIS PROFILES-- |              |                  |               |             |             |                         |
|------------------|--------------|------------------|---------------|-------------|-------------|-------------------------|
| Index            | Profile Name | Refresh Timer(s) | AIS RDI State | No: Of Lsps | Mode        | Service Degrade Support |
| 1                | Default      | 10               | Enabled       | 1           | Single Byte | Disabled                |
| 2                | ais1         | 10               | Enabled       | 0           | Single Byte | Disabled                |

### To delete an AIS/LDI profile

#### 4 Delete a specific AIS/LDI profile:

```
ais profile delete profile <profile-name>
```

where

<profile-name> Specifies the name of the profile for to delete.

**Note:** A profile cannot be deleted when it is in use by any LSP.

For example:

## 18-8 Alarm Indication Signal with Link Down Indication

---

```
ais profile delete profile ais1
```

Verify that the AIS/LDI profile is deleted:

```
ais profile show
```

| AIS PROFILES |              |                     |                  |             |             |                            |
|--------------|--------------|---------------------|------------------|-------------|-------------|----------------------------|
| Index        | Profile Name | Refresh<br>Timer(s) | AIS RDI<br>State | No: Of Lsps | Mode        | Service Degrade<br>Support |
| 1            | Default      | 10                  | Enabled          | 1           | Single Byte | Disabled                   |

—end—

## Procedure 18-5

# Configuring AIS/LDI on a static ingress associated TP tunnel (on an LER)

A static ingress associated TP tunnel comprises:

- one static ingress unidirectional TP tunnel
- one static egress unidirectional TP tunnel

You can

- configure AIS/LDI when you create a static ingress associated TP tunnel
- configure or change AIS/LDI on an existing static ingress associated TP tunnel

**Note:** If AIS/LDI is globally enabled but not configured for the tunnel, SNMP traps are generated when an AIS/LDI message is received.

| Step | Action |
|------|--------|
|------|--------|

### To create a static ingress associated TP tunnel with AIS/LDI configured

- 1 Create a static ingress associated TP tunnel with AIS/LDI configured:

```
gmpls tp-tunnel create static-ingress-assoc <static-
ingress-assoc>
{forward-tunnel <MPLS ingress tp-unidir-tunnel>}
{reverse-static-tunnel <MPLS static Egress unidir tp-
tunnel>}
[bfd-monitor <enable | disable>]
[bfd-profile <MPLS BFD Profile List>]
[ais-monitor <enable | disable>]
[ais-profile <AIS Profile List>]
```

For example, create static ingress associated TP tunnel “sAs-198-152” with AIS/LDI configured:

```
gmpls tp-tunnel create static-ingress-assoc sAs-198-152 forward-tunnel sAs-
198-165-159-166-152 reverse-static-tunnel sAs-152-236-198 bfd-monitor enable
ais-monitor enable
```

(Refer to [“5150-198 LER configuration” on page 18-20](#) for an example of this command in context of the [“Example topology: AIS over static MPLS-TP associated tunnels” on page 18-19](#) illustrated in [“Static MPLS-TP associated tunnels” on page 18-19](#).)

- 2 Enable the static ingress associated TP tunnel:

```
gmpls tp-tunnel enable static-ingress-assoc <static-
ingress-assoc>
```

- 3 Verify that the static ingress associated tunnel was created:  

```
gmpls tp-tunnel show static-ingress-assoc <static-ingress-assoc>
```

***To configure or modify AIS/LDI on an existing static ingress associated TP tunnel***

- 4 Configure AIS/LDI on an existing static static-ingress-assoc TP tunnel:  

```
gmpls tp-tunnel set static-ingress-assoc <static-ingress-assoc>
[bfd-monitor <enable | disable>]
[bfd-profile <MPLS BFD Profile List>]
[ais-monitor <enable | disable>]
[ais-profile <AIS Profile List>]
```
- 5 Verify that AIS/LDI was configured or modified correctly:  

```
gmpls tp-tunnel show static-ingress-assoc <static-ingress-assoc>
```

—end—

## Procedure 18-6

# Configuring AIS/LDI on a static transit unidirectional TP tunnel (on an LSR)

You can

- configure AIS/LDI when you create a static transit unidirectional TP tunnel
- configure or change AIS/LDI on an existing static transit unidirectional TP tunnel

**Note:** If AIS/LDI is globally enabled but not configured for the tunnel, SNMP traps are generated when an AIS/LDI message is received.

| Step | Action |
|------|--------|
|------|--------|

### To create a static transit unidirectional TP tunnel with AIS/LDI configured

- 1 Create a static transit unidirectional TP-tunnel with AIS/LDI configured:

```
gmpls tp-tunnel create static-transit-unidir <static-transit-unidir> {dest-ip <IP address>} {src-ip <IP address>} {next-hop-ip <IP address>} {forward-out-label <NUMBER: 16-1044479>} {forward-in-label <NUMBER: 16-1044479>} [cos-profile <MPLS Tunnel COS Profile>] {ttl-policy <fixed | decrement>} [fixed-ttl <NUMBER: 1-255>] [prev-hop-ip <IP address>] [ais-monitor <enable | disable>] [ais-profile <AIS Profile List>]
```

For example, create static transit unidirectional TP-tunnel “Transit1-1” with AIS/LDI configured:

```
gmpls tp-tunnel create static-transit-unidir Transit1-1 dest-ip 4.4.4.4 src-ip 1.1.1.1 next-hop-ip 30.1.1.2 forward-out-label 4101 forward-in-label 101 prev-hop-ip 10.1.1.1 ais-monitor enable ais-profile dp-ais-profile-1
```

- 2 Enable the dynamic unidirectional TP tunnel:

```
gmpls tp-tunnel enable static-transit-unidir <static-transit-unidir>
```

### To configure or modify AIS/LDI on an existing static transit unidirectional TP tunnel

- 3 Configure AIS/LDI on an existing static transit unidirectional TP tunnel:

```
gmpls tp-tunnel set static-transit-unidir <static-transit-unidir> [ais-monitor <enable | disable>] [ais-profile <AIS Profile List>]
```

- 4 Verify that AIS/LDI was configured or modified correctly:

```
gmpls tp-tunnel show static-transit-unidir <static-transit-unidir>
```

—end—

## Procedure 18-7

### Configuring AIS/LDI on a static ingress co-routed TP tunnel (on an LER)

You can

- configure AIS/LDI when you create a static ingress co-routed TP tunnel
- configure or change AIS/LDI on an existing static ingress co-routed TP tunnel

**Note:** If AIS/LDI is globally enabled but not configured for the tunnel, SNMP traps are generated when an AIS/LDI message is received.

| Step | Action |
|------|--------|
|------|--------|

#### *To create a static ingress co-routed TP tunnel with AIS/LDI configured*

- 1 Create a static ingress co-routed TP tunnel with AIS/LDI configured:  

```
gmpls tp-tunnel create static-ingress-corout <static-ingress-corout> {dest-ip <IP address>} {next-hop-ip <IP address>} {forward-out-label <NUMBER: 16-1044479>} {reverse-in-label <NUMBER: 16-1044479>} [cos-profile <MPLS Tunnel COS Profile>] {ttl-policy <fixed>} [fixed-ttl <NUMBER: 1-255>] [backup-tunnel <MPLS ingress primary tp corout tunnel>] [reversion-hold-time <NUMBER: 0-3600>] [tunnel-reversion <on | off>] [bfd-monitor <enable | disable>] [bfd-profile <MPLS BFD Profile List>] [ais-monitor <enable | disable>] [ais-profile <AIS Profile List>]
```
- 2 Enable the static ingress co-routed TP tunnel:  

```
gmpls tp-tunnel enable static-ingress-corout <static-ingress-corout>
```
- 3 Verify that the static ingress tunnel was created:  

```
gmpls tp-tunnel show static-ingress-corout <static-ingress-corout>
```

#### *To configure or modify AIS/LDI on an existing static ingress co-routed TP tunnel*

- 4 Configure AIS/LDI on an existing static ingress co-routed TP tunnel:  

```
gmpls tp-tunnel set static-ingress-corout <static-ingress-corout> [ais-monitor <enable | disable>] [ais-profile <AIS Profile List>]
```
- 5 Verify that AIS/LDI was configured or modified correctly:  

```
gmpls tp-tunnel show static-ingress-corout <static-ingress-corout>
```

—end—

## Procedure 18-8

# Configuring AIS/LDI on a static egress co-routed TP tunnel (on an LER)

You can

- configure AIS/LDI when you create a static egress co-routed TP tunnel
- configure or change AIS/LDI on an existing static egress co-routed TP tunnel

**Note:** If AIS/LDI is globally enabled but not configured for the tunnel, SNMP traps are generated when an AIS/LDI message is received.

| Step | Action |
|------|--------|
|------|--------|

### To create a static egress co-routed TP tunnel with AIS/LDI configured

- 1 Create a static egress co-routed TP tunnel with AIS/LDI configured:  

```
gmpls tp-tunnel create static-egress-corout <static-egress-corout> src-ip <IP address> prev-hop-ip <IP address> forward-in-label <NUMBER: 16-1044479> reverse-out-label <NUMBER: 16-1044479> [cos-profile <MPLS Tunnel COS Profile>] [ttl-policy <fixed>] [fixed-ttl <NUMBER: 1-255>] [reversion-hold-time <NUMBER: 0-3600>] [tunnel-reversion <on | off>] [backup-tunnel <MPLS static egress primary tp-tunnel>] [bfd-monitor <enable | disable>] [bfd-profile <MPLS BFD Profile List>] [ais-monitor <enable | disable>] [ais-profile <AIS Profile List>]
```
- 2 Enable the static egress co-routed TP tunnel:  

```
gmpls tp-tunnel enable static-egress-corout <static-egress-corout>
```
- 3 Verify that the static ingress tunnel was created:  

```
gmpls tp-tunnel show static-egress-corout <static-egress-corout>
```

### To configure or modify AIS/LDI on an existing static egress co-routed TP tunnel

- 4 Configure AIS/LDI on an existing static egress co-routed TP tunnel:  

```
gmpls tp-tunnel set static-egress-corout <static-egress-corout> [ais-monitor <enable | disable>] [ais-profile <AIS Profile List>]
```
- 5 Verify that AIS/LDI was configured or modified correctly:  

```
gmpls tp-tunnel show static-egress-corout <static-egress-corout>
```

—end—



## Procedure 18-9

### Configuring AIS/LDI on a static transit co-routed TP tunnel (on an LSR)

You can

- configure AIS/LDI when you create a static transit co-routed TP tunnel
- configure or change AIS/LDI on an existing static transit co-routed TP tunnel

**Note:** If AIS/LDI is globally enabled but not configured for the tunnel, SNMP traps are generated when an AIS/LDI message is received.

| Step | Action |
|------|--------|
|------|--------|

#### To create a static transit co-routed TP tunnel with AIS/LDI configured

- 1 Create a static transit co-routed TP tunnel with AIS/LDI configured:

```
gmpls tp-tunnel create
static-transit-corout <static-transit-corout>
[logical-id <NUMBER>]
{src-ip <IP address>}
{forward-out-label <NUMBER: 16-1044479>}
{forward-in-label <NUMBER: 16-1044479>}
{reverse-out-label <NUMBER: 16-1044479>}
{reverse-in-label <NUMBER: 16-1044479>}
[cos-profile <MPLS Tunnel COS Profile>]
{ttl-policy <fixed | decrement>}
[fixed-ttl <NUMBER: 1-255>]
[ais-monitor <enable | disable>]
[ais-profile <AIS Profile List>]
```

For example, create static transit co-routed “sCo-198-165-159-166-152” with AIS/LDI configured:

```
gmpls tp-tunnel create static-transit-corout sCo-198-165-159-166-152 dest-ip
7.152.10.10 src-ip 7.198.10.10 next-hop-ip 203.0.113.1 prev-hop-ip
198.51.100.0 forward-out-label 26 forward-in-label 25 reverse-out-label 25
reverse-in-label 26 ais-monitor enable
```

(Refer to [“3930-165 LSR configuration:” on page 18-23](#), [“3916-159 LSR configuration:” on page 18-23](#) and [“3931-166 LSR configuration:” on page 18-23](#) for an example of this command in context of the [“Example topology: AIS over static MPLS-TP co-routed tunnels” on page 18-22](#)” illustrated in [“Static MPLS-TP co-routed tunnels” on page 18-22](#).)

- 2 Enable the static transit co-routed TP tunnel:

```
gmpls tp-tunnel enable static-transit-corout <static-
transit-corout>
```

- 3 Verify that the static transit tunnel was created:

```
gmpls tp-tunnel show static-transit-corout <static-
transit-corout>
```

***To configure or modify AIS/LDI on an existing static transit co-routed TP tunnel***

- 4 Configure AIS/LDI on an existing static transit co-routed TP tunnel:

```
gmpls tp-tunnel set static-transit-corout <static-
transit-corout>
forward-out-label <NUMBER: 16-1044479>]
[forward-in-label <NUMBER: 16-1044479>]
[reverse-out-label <NUMBER: 16-1044479>]
[reverse-in-label <NUMBER: 16-1044479>]
[cos-profile <MPLS Tunnel COS Profile>]
{ttl-policy <fixed | decrement>}
[fixed-ttl <NUMBER: 1-255>]
[ais-monitor <enable | disable>]
[ais-profile <AIS Profile List>]
```

- 5 Verify that AIS/LDI was configured or modified correctly:

```
gmpls tp-tunnel show static-egress-corout <static-egress-
corout>
```

—end—

## Procedure 18-10

### Displaying AIS/LDI global statistics

You can display AIS/LDI global statistics.

**Note:** The column 'Previous Hop Fault Count' lists the total number of times the previous hop has gone down, whether or not the tunnel itself has seen all of those instances of the previous hop going down. The previous hop down count is kept separate from the tunnel and so reflects the actual count, not what the tunnel has seen.

| Step | Action                                                                                                                    |
|------|---------------------------------------------------------------------------------------------------------------------------|
| 1    | Display AIS/LDI global LER and LSR statistics.<br><pre>ais statistics show</pre> <p style="text-align: center;">—end—</p> |

### Example

The following example shows sample output for an LER.

```
3928*> ais statistics show
```

| --AIS LER TABLE-- |              |                        |                 |                 |                 |          |          |
|-------------------|--------------|------------------------|-----------------|-----------------|-----------------|----------|----------|
| Tunnel Name       | Previous Hop | Previous Hop Interface | Fault Set Count | LSP Fault state | LSP Fault Count | SD state | SD Count |
| cr1               | 145.1.1.2    | unnum1                 | 2               | Yes             | 2               | No       | 22       |

| --AIS LSR TABLE-- |              |                        |                 |                 |          |          |  |
|-------------------|--------------|------------------------|-----------------|-----------------|----------|----------|--|
| Tunnel Name       | Previous Hop | Previous Hop Interface | Fault Set Count | LSP Fault State | SD Count | SD state |  |

The tunnel named `cr1` is a static ingress associated tunnel. The output for this tunnel in the AIS LER TABLE above shows that:

- The previous hop (upstream neighbor) IP address is 145.1.1.2.
- The name of the previous hop interface is unnum1.
- The fault set fault count is two (two faults were detected).
- The LSP fault state is yes.
- The LSP fault count is two (two faults were detected on the LSP).
- Service degrade detection is disabled.
- There were 22 service degrade detections.

The following example shows sample output for an LSR.

```
> ais statistics show
```

| +--AIS LSR TABLE--+-- |              |                        |                 |                 |          |          |
|-----------------------|--------------|------------------------|-----------------|-----------------|----------|----------|
| Tunnel Name           | Previous Hop | Previous Hop Interface | Fault Set Count | LSP Fault State | SD Count | SD state |
| crl                   | 135.1.1.1    | unnum2                 | 29              | No              | 0        | No       |
| crl                   | 145.1.1.1    | unnum3                 | 5               | Yes             | 0        | No       |

The tunnel named staticTransit2 is a static transit tunnel. The output for this tunnel in the AIS LSR TABLE above shows that:

- The previous hop (upstream neighbor) IP address is 10.10.11.1.
- The names of the previous hop interfaces are unnum2 and unnum3.
- There are two previous hop faults.
- As a result of those two hop faults, 53 AIS messages were transmitted downstream towards the LER.

## Procedure 18-11

### Clearing AIS/LDI global statistics

You can reset AIS/LDI global statistics to zero.

| Step | Action |
|------|--------|
|------|--------|

- 1 If you belong to the admin or superuser user group, clear AIS/LDI global statistics:  

```
ais statistics clear
```

After statistics are cleared the “ais statistics show” command displays a table similar to the example below:

```
ais statistics show
```

| AIS LER TABLE |              |                        |                 |                 |                 |          |          |
|---------------|--------------|------------------------|-----------------|-----------------|-----------------|----------|----------|
| Tunnel Name   | Previous Hop | Previous Hop Interface | Fault Set Count | LSP Fault state | LSP Fault Count | SD state | SD Count |
| crl           | 135.1.1.2    |                        | 0               | Yes             | 0               | No       | 0        |

| AIS LSR TABLE |              |                        |                 |                 |          |          |
|---------------|--------------|------------------------|-----------------|-----------------|----------|----------|
| Tunnel Name   | Previous Hop | Previous Hop Interface | Fault Set Count | LSP Fault State | SD Count | SD state |
|               |              |                        |                 |                 |          |          |

—end—



### **5150-198 LER configuration**

```
gmpls tp-tunnel create static-ingress-unidir sAs-198-165-159-166-152 dest-ip
7.152.10.10 next-hop-ip 198.51.100.1 forward-out-label 16

gmpls tp-tunnel create static-ingress-unidir sAs-198-152-bkp dest-ip
7.152.10.10 next-hop-ip 7.30.30.31 forward-out-label 32 backup-tunnel sAs-
198-165-159-166-152

gmpls tp-tunnel create static-egress-unidir sAs-152-236-198 src-ip
7.152.10.10 forward-in-label 21 prev-hop-ip 7.70.70.71

gmpls tp-tunnel create static-egress-unidir sAs-152-198-bkp src-ip
7.152.10.10 forward-in-label 32 prev-hop-ip 7.30.30.31

gmpls tp-tunnel create static-ingress-assoc sAs-198-152 forward-tunnel sAs-
198-165-159-166-152 reverse-static-tunnel sAs-152-236-198 bfd-monitor enable
ais-monitor enable

gmpls tp-tunnel create static-ingress-assoc sAs-198-152-bkp forward-tunnel
sAs-198-152-bkp reverse-static-tunnel sAs-152-198-bkp bfd-monitor enable ais-
monitor enable
```

### **3930-165 LSR configuration**

```
gmpls tp-tunnel create static-transit-unidir sAs-198-165-159-166-152 dest-ip
7.152.10.10 src-ip 7.198.10.10 next-hop-ip 203.0.113.1 forward-out-label 17
forward-in-label 16 prev-hop-ip 198.51.100.0 ais-monitor enable
```

### **3916-159 LSR configuration**

```
gmpls tp-tunnel create static-transit-unidir sAs-198-165-159-166-152 dest-ip
7.152.10.10 src-ip 7.198.10.10 next-hop-ip 7.60.60.60 forward-out-label 18
forward-in-label 17 prev-hop-ip 203.0.113.0 ais-monitor enable
```

### **3931-166 LSR configuration**

```
gmpls tp-tunnel create static-transit-unidir sAs-198-165-159-166-152 dest-ip
7.152.10.10 src-ip 7.198.10.10 next-hop-ip 192.0.2.0 forward-out-label 19
forward-in-label 18 prev-hop-ip 7.60.60.61 ais-monitor enable
```

### **3916-236 LSR configuration**

```
gmpls tp-tunnel create static-transit-unidir sAs-152-236-198 dest-ip
7.198.10.10 src-ip 7.152.10.10 next-hop-ip 7.70.70.70 forward-out-label 21
forward-in-label 20 prev-hop-ip 7.80.80.80 ais-monitor enable
```



## Procedure 18-13


### Example topology: AIS over static MPLS-TP co-routed tunnels

Figure 18-2 illustrates an AIS configuration for associated primary and co-routed backup tunnel configuration.

Configuration for each node in that topology is detailed in the following sections:

- “5150-198 LER configuration” on page 18-20
- “3930-165 LSR configuration:” on page 18-23
- “3916-159 LSR configuration:” on page 18-23
- “3931-166 LSR configuration:” on page 18-23

**Figure 18-2**  
Static MPLS-TP co-routed tunnels



**5150-198 LER configuration:**

```
gmpls tp-tunnel create static-ingress-corout sCo-198-165-159-166-152 dest-ip
7.152.10.10 next-hop-ip 198.51.100.1 forward-out-label 25 reverse-in-label
25 bfd-monitor enable ais-monitor enable
```

```
gmpls tp-tunnel create static-ingress-corout sCo-198-152-bkp dest-ip
7.152.10.10 next-hop-ip 7.30.30.31 forward-out-label 30 reverse-in-label 30
backup-tunnel sCo-198-165-159-166-152 bfd-monitor enable ais-monitor enable
```

**3930-165 LSR configuration:**

```
gmpls tp-tunnel create static-transit-corout sCo-198-165-159-166-152 dest-ip
7.152.10.10 src-ip 7.198.10.10 next-hop-ip 203.0.113.1 prev-hop-ip
198.51.100.0 forward-out-label 26 forward-in-label 25 reverse-out-label 25
reverse-in-label 26 ais-monitor enable
```

**3916-159 LSR configuration:**

```
gmpls tp-tunnel create static-transit-corout sCo-198-165-159-166-152 dest-ip
7.152.10.10 src-ip 7.198.10.10 next-hop-ip 7.60.60.60 prev-hop-ip 203.0.113.0
forward-out-label 27 forward-in-label 26 reverse-out-label 26 reverse-in-
label 27 ais-monitor enable
```

**3931-166 LSR configuration:**

```
gmpls tp-tunnel create static-transit-corout sCo-198-165-159-166-152 dest-ip
7.152.10.10 src-ip 7.198.10.10 next-hop-ip 192.0.2.0 prev-hop-ip 7.60.60.61
forward-out-label 28 forward-in-label 27 reverse-out-label 27 reverse-in-
label 28 ais-monitor enable
```



---

## Virtual private LAN service integrated routing and bridging

---

Virtual Private LAN Service (VPLS) provides L2 Ethernet-based multipoint-to-multipoint communication over MPLS networks. Geographically dispersed sites can share an Ethernet broadcast domain by connecting sites through pseudowires (PW). VPLS extends a LAN to the edge of the provider network and then the provider network emulates a switch or bridge to connect all the LANs to create a single bridged LAN.

Integrated Routing and Bridging (IRB) allows a router to perform bridging and/or routing on the same interface based on the destination MAC (DMAC) address. If the DMAC address of the Ethernet frame is not addressed to the router, the frame is L2 switched/bridged. If the destination MAC address of the Ethernet frame is addressed to the router, the frame is L2 terminated and IP forwarded if the payload is an IP packet.

For IRB over VPLS, the provider edge designates an entire virtual-switch instance as a “virtual” IP interface to perform IRB functionality regardless of whether the Ethernet frame is received on a member attachment circuit (AC) or PW.

**Note:** MPLS VPLS IRB is supported on the 3930, 3932 and 5160 devices.

### VPLS IRB virtual switch configuration

An MPLS VPLS IRB virtual switch requires an Ethernet virtual switch with the mode “vpls” and reserved VLAN to be configured. The virtual switch is then associated with an IP interface to become a VPLS IRB virtual switch. All virtual switches with an IP interface are part of the same routing domain. The reserved VLAN ID must be unique across the system, and it cannot overlap with any other VLAN which is already associated with another IP interface or VPLS IRB virtual switch.

The following restrictions apply:

- All member ACs belonging to the same VPLS IRB instance must be configured with the same VLAN encapsulation. All VPLS IRB instances across the network must be configured with the same VLAN encapsulation if they are associated with the same IP subnet.

The existing VLAN transformation functionality can be used, but if it is not configured correctly, problems can result. For example, the VLAN encapsulation for the IP packet is not consistent within and/or among local and remote VPLS IRB instances associated with the same IP subnet. As a result, the existing VLAN transformation functionality is not supported for the VPLS IRB virtual switch.

- The untagged member AC is not supported.
- Double tagged (QinQ) VLAN encapsulation is not supported.

The following functionality is disabled on an VPLS IRB IP interface and a VPLS switch:

- RSVP-TE interface configuration is blocked
- LDP interface is disabled.
- Association of remote management interface towards VPLS IRB virtual switch is disabled.
- ACL configuration on a VPLS IRB IP interface is blocked.
- MPLS static and/or dynamic TNL configuration on a next-hop IP address associated with VPLS IRB IP interface is blocked.

### VPLS IRB IP interface

A virtual IP interface must be created and associated with the VPLS IRB VS.

The virtual IP interface's operational state and characteristic information is similar to a typical VLAN-based IP interface. In certain cases, this information is determined by the member AC and PW state and characteristic information of the corresponding VPLS IRB instance VS:

- The MTU must be smaller than the corresponding NNI port max-frame size (considering the additional overhead introduced by PW/TNL and the outer Ethernet header).
- When the first AC or PW is operationally up, the IP interface is operationally up.
- When all ACs and PWs are operationally down, the IP interface is operationally down.

Otherwise, the VPLS IRB IP interface can be used as the normal VLAN-based IP interface where various existing IP functionality is supported and capable of IP traffic hardware forwarding:

- IPv4 static routing

- IPv4 dynamic routing protocols (OSPF)
- IP OAM tools
  - Ping and traceroute
  - In-band management
    - Telnet
    - SSH
    - xFTP
    - SNMP
    - Radius
    - TACACS
    - RMON
    - NTP
    - DNS
    - Syslog

The IP interface can also be used as a normal IP interface for the following configurations:

- static ARP
- OSPF routing

For OSPF routing, the OSPF area route filtering capability allows or denies routes leaked between areas. The OSPF passive interface capability simplifies the distribution of the stub network route without affecting the neighbor relationship. This behavior is not unique to VPLS IRB.

To route or forward IP packets multiple hops away, you must either configure static routes (using the `ip route` command) to explicitly configure the routes or enable OSPF to let the protocol distribute and install the routes. In both cases, the route eventually makes it into the IP forwarding table (FIB) and is pushed down into the hardware to perform IP HW forwarding. VPLS IRB does not change the existing IP routing and OSPF configuration.

Deleting or disabling an IP interface is unchanged with VPLS IRB. The mode of a newly created virtual switch cannot be changed. When an IP interface is disabled, IP forwarding over the IP interface is also disabled regardless of whether the IP interface is associated with the VPLS IRB virtual switch.

VPLS IRB supports setting all other optional attributes previously supported by existing-based interface with these exceptions:

- Unique per IP interface MAC address
- unnumbered IP interface

- Various software and hardware-based access control lists

## Unicast IPv4 hardware forwarding over a VPLS IRB IP interface

Table 19-1 summarizes the maximum bidirectional IPv4 unicast hardware (HW) forwarding throughput supported based on additional passes required on internal/external loopback port(s). This assumes 10 Gbps of internal/external loopback port(s) are used.

**Table 19-1**  
**Unicast IPv4 HW forwarding traffic scenarios**

| Scenario | Between       | 3930/3932      | 5160           | Notes                                               |
|----------|---------------|----------------|----------------|-----------------------------------------------------|
| 1        | AC - AC       | 5.0 Gbps       | Not applicable | Loopback port not used for 5160                     |
| 2        | AC - PW       | 2.5 Gbps       | 5.0 Gbps       |                                                     |
| 3        | IP IF - AC    | 10.0 Gbps      | Not applicable | Loopback port not used for 5160                     |
| 4        | PW - PW       | 1.67 Gbps      | 2.5 Gbps       |                                                     |
| 5        | IP IF - PW    | 3.33 Gbps      | 5.0 Gbps       |                                                     |
| 6        | IP IF - IP IF | Not applicable | Not applicable | Loopback port not used for both 3930, 3932 and 5160 |

## IPv4 local termination and origination over a VPLS IRB IP interface

IPv4 local termination and origination is supported over the VPLS IRB IP interface. VPLS IRB supports IPv4 unicast/broadcast/multicast IPv4 local termination for packets received on a member

- UNI AC of the VPLS IRB instance
- NNI PW of the VPLS IRB instance
- UNI AC of the VPLS IRB instance
- NNI PW of the VPLS IRB instance

IPv4 local origination and termination packet is encapsulated as Ethernet L2 VPN PW payload when it is sent or received over the NNI PW.

The following IPv4 packet local termination and origination is supported:

- ICMP (for ping and trace route)
- OSPF
- TCP
- UDP

## Broadcast/multicast L2 local termination and origination over a VPLS IRB IP interface

Broadcast/multicast L2 local termination and origination over the VPLS IRB IP interface is supported. Broadcast/multicast L2 local termination for packets received on member

- UNI AC of the VPLS IRB instance
- NNI PW of the VPLS IRB instance

Broadcast/multicast L2 local origination for packets transmitted on member

- UNI AC of the VPLS IRB instance
- NNI PW of the VPLS IRB instance

Broadcast/multicast L2 local termination and origination packet is encapsulated as Ethernet L2 VPN PW payload when it is sent or received over the NNI PW.

The following types of broadcast/multicast L2 frame local termination and origination is supported:

- ARP (L2 broadcast)
- OSPF (L2 multicast)
- IS-IS (L2 multicast)

## VPLS IRB IP interface supported functionality

VPLS IRB IP interface supports the following functionality:

- DiffServ based Per-Hop-Behavior (PHB) over the VPLS IRB IP interface
- Ping and Traceroute
- OSPF
- ARP
- L2 bridging and L3 IP forwarding simultaneously

## IRB specific IFP resources management

The resource-managed feature “advanced-I3” manages IRB specific resources.

On the 3930 and 3932, it is likely that all field processor resources have been allocated, and another feature’s resource allocation must be reduced first. The 5160 device may, by default, have some faceplate resources that are not allocated. If this is the case, then another feature’s allocation does not need to be reduced. Otherwise, another feature’s allocation must be reduced for VPLS IRB.



**Note:** Due to limited hardware resources available on various platforms, VPLS IRB must be enabled or disabled to reserve or un-reserve the system resources required to use MPLS VPLS IRB.

## Scalability

“[Scalability requirements](#)” on page 19-6 describes the scalability requirements for MPLS VPLS IRB.

**Table 19-2**  
**Scalability requirements**

| Device              | Support                                                                              |
|---------------------|--------------------------------------------------------------------------------------|
| 3930 and 3932       | 32 IP interfaces<br>• 16 VPLS IRB virtual switches<br>• 16 VPLS IRB IP interfaces    |
| 5160                | 256 IP interfaces<br>• 128 VPLS IRB virtual switches<br>• 128 VPLS IRB IP interfaces |
| 3930, 3932 and 5160 | 4K routing table entries                                                             |
|                     | 4K IP hardware forwarding table entries                                              |
|                     | 1K next-hop entries                                                                  |
|                     | 1K IP hardware adjacency table entries                                               |

## Feature considerations

The VPLS IRB feature has the following feature considerations:

- IRB functionality is not supported with Control Word (CW).
- Limitations require IP packets to pass through the hardware pipeline multiple times in specific traffic forwarding scenarios by using internal or external ports reserved for the MPLS VPLS IRB feature.
- The IPv4 hardware forwarding throughput is capped by the available bandwidth of the internal or external ports for specific traffic forwarding scenarios.
- External face plate ports must be used if there are no hidden internal ports with sufficient bandwidth.
- Using a LAG for multi-pass purposes to increase the IP hardware forwarding bandwidth it is not supported.

- All UNI ACs attached to the VPLS IRB virtual switch belonging to the same IP subnet across the network need to use the same VLAN encapsulation (same VLAN ID).
- The untagged member AC is not supported by the MPLS VPLS IRB VS IP interface.
- QinQ (double-tagged VLAN) is not supported by MPLS VPLS IRB VS and/or an associated IP interface. Double-tagged VLAN in the regular VLAN-based IP interface is not supported.
- Only IPv4 IP hardware forwarding is supported over MPLS VPLS IRB and the associated IP interface. IPv6 is not supported by MPLS VPLS IRB. IPv6 hardware forwarding over the regular VLAN-based IP interface is not supported.
- Only one IP interface/subnet can be associated with a single VPLS IRB VS. Associating multiple IP interfaces/subnets to the same VPLS IRB VS is not supported.
- There is no separate VRF support for the VPLS IRB VS and its associated IP interface. All the IP interface (regular VLAN based or VPLS IRB VS based) will reside in the same default VRF (global VRF).
- OSPF is supported as a dynamic routing protocol for MPLS VLS IRB.
- The MPLS VPLS IRB feature is supported on the 3930, 3932 and the 5160.
- The default RCOS, FCOS and RCOS to queue mapping is supported for MPLS VPLS IRB.
- The IRB IP interface cannot be excluded from the CSPF tunnel path calculation even if RSVP-TE is not enabled on the IRB IP interface. This means that the CSPF tunnel path calculation still includes the IRB IP interface. The IRB IP interface cannot co-exist with a dynamically signaled tunnel because the IRB IP interface might be unintentionally considered as part of the CSPF tunnel path calculation.
- Traffic profiling on the VPLS IRB VS is not supported due to the following reasons:
  - The internal AC over the loopback port(s) are accounted for which is not intended.
  - For traffic dropping, IRB is not able to differentiate control and data traffic. This can cause OSPF and IP BFD adjacencies to drop which will lead to route withdrawal and complete dropping of traffic.
- Mesh pseudowire is not supported due to the limitation of the OSPF broadcast network type and DR/BDR election. The suitable OSPF network type is NBMA, which is not currently supported by SAOS.

- In the dual-homing scenario where MTUs are connected to two AggrGWs, every pseudowire must be attached to their own independent VPLS IRB virtual switch.
- The PCP value of the VLAN received on the terminal side is determined by the DSCP value of the IP packet sent by the sender in IRB. This is because IRB resolves the internal RCoS value by the DSCP mapping in the RCoS map regardless of what the access port is. The AC port's resolved CoS policy or encap CoS policy is set in the virtual switch.

If the access port's (AC port) resolve CoS policy is default (dot1dag), then the queuing will be different in the node because the packet is looped internally. One iteration is queued based on the resolve CoS policy of the port or VS (if vs-inherit), and the second iteration is queued based on the DSCP mapping RCoS value.

To achieve consistent queuing for both iterations, there are two options:

- Set the port's resolve CoS policy to dscp mapped, and VS encap cos policy to port-inherit.
- Set the sent IP traffic's VLAN PCP value to match the DSCP RCoS value.

See the following procedures:

- [“Reserving the resource manager advanced-l3 feature resource for the 3930, 3932 and the 5160” on page 19-10](#)
- [“Configuring the port to the list of system reserved ports for VPLS IRB” on page 19-11](#)
- [“Enabling or disabling VPLS IRB” on page 19-12](#)
- [“Clearing VPLS IRB statistics” on page 19-13](#)
- [“Configuring an MPLS VPLS IRB virtual switch” on page 19-14](#)
- [“Configuring an MPLS VPLS IP interface associated with the VPLS IRB virtual switch” on page 19-15](#)
- [“Configuring static ARP over the virtual IP interface” on page 19-16](#)
- [“Configuring OSPF passive interface on the MTU for IP interfaces connecting towards eNBs” on page 19-17](#)
- [“Configuring OSPF area route filtering to separate OSPF route distribution” on page 19-18](#)
- [“Displaying system reserved ports for VPLS IRB” on page 19-19](#)
- [“Displaying the status of MPLS VPLS IRB” on page 19-20](#)
- [“Displaying VPLS IRB bandwidth usage statistics” on page 19-21](#)
- [“Displaying an MPLS VPLS IRB virtual switch” on page 19-22](#)
- [“Displaying MPLS VPLS IRB IP interface information” on page 19-23](#)

- “Displaying the static ARP entry over the VPLS IRB IP interface” on page 19-26
- “Configuration example using integrated routing and bridging and MTUs” on page 19-27
- “Configuration example using MTUs alternative using IP interface” on page 19-30
- “Configuration example using aggregate gateway (5160)” on page 19-34
- “Configuration example using aggregate gateway (5160)” on page 19-37

## Procedure 19-1

### Reserving the resource manager advanced-l3 feature resource for the 3930, 3932 and the 5160

---

If required you can reduce another feature's resource allocation on the 3930, 3932 or 5160, to use VPLS IRB.

**Note:** You must enable/disable the VPLS IRB feature to reserve/un-reserve the system resources required to use VPLS IRB.

| Step | Action                                                                                                                                                                       |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Configure the resource-manager advanced-l3 resource classifier count:<br><code>resource-manager pool set feature advanced-l3 resource classifier count &lt;NUMBER&gt;</code> |
| 2    | Configure the resource-manager advanced-l3 resource:<br><code>resource-manager pool set feature advanced-l3 resource counter count &lt;NUMBER&gt;</code>                     |
| 3    | Validate the configuration:<br><code>resource-manager validate</code>                                                                                                        |
| 4    | Save the configuration:<br><code>config save</code>                                                                                                                          |
| 5    | Reboot the device:<br><code>chassis reboot</code>                                                                                                                            |

—end—

### Example

The following example reserves the resource manager advanced l-3 feature resource for the 3930 and 3932:

```
> resource-manager pool set feature advanced-l3 resource classifier count 256
> resource-manager pool set feature advanced-l3 resource counter count 256
> resource-manager validate
> config save
> chassis reboot
```

The following example reserves the resource manager advanced l-3 feature resource for the 5160:

```
> resource-manager pool set feature advanced-l3 resource classifier count 512
> resource-manager pool set feature advanced-l3 resource counter count 512
> resource-manager validate
> config save
> chassis reboot
```

## Procedure 19-2

### Configuring the port to the list of system reserved ports for VPLS IRB

You can add the specified port to the list of system reserved ports to use for VPLS IRB. You can also delete the port from the list of system reserved ports.

| Step | Action |
|------|--------|
|------|--------|

#### *To add the port to the list of system reserved ports for VPLS IRB*

- 1 Add the port to the list of system reserved ports for VPLS IRB:  

```
chassis system-rsvd-port add port <port> feature vpls-irb
```

where

port <port> allows the external faceplate port to be reserved for VPLS IRB.

#### *To remove the port from the list of system reserved ports for VPLS IRB*

- 2 Remove the port to the list of system reserved ports for VPLS IRB:  

```
chassis system-rsvd-port remove port <port>
```

—end—

## Procedure 19-3

### Enabling or disabling VPLS IRB

---

Due to limited hardware resources, you must enable or disable VPLS IRB to reserve or un-reserve system resources required to use MPLS VPLS IRB.

**Note:** If all the required system resources are not available for VPLS IRB, the enable attempt will be rejected. VPLS IRB cannot be disabled until all the IP interfaces associated with VPLS VS are removed.

---

| Step | Action |
|------|--------|
|------|--------|

---

- |   |                                                                                                     |
|---|-----------------------------------------------------------------------------------------------------|
| 1 | Enable or disable VPLS IRB:<br><code>vpls-irb [&lt;enable disable&gt;]</code><br><code>—end—</code> |
|---|-----------------------------------------------------------------------------------------------------|

## Procedure 19-4

### Clearing VPLS IRB statistics

Statistics can be cleared for the internal AC and IP interface port on a 3930 and 3932, or if needed for the reserved loopback port on the 5150.

| Step | Action |
|------|--------|
|------|--------|

- |   |                                            |
|---|--------------------------------------------|
| 1 | Clear VPLS IRB bandwidth usage statistics: |
|---|--------------------------------------------|

```
vpls-irb clear [statistics]
```

where

clear [statistics} displays information and loopback sports statistics (optional).

—end—



## Procedure 19-5

### Configuring an MPLS VPLS IRB virtual switch

To configure an MPLS VPS IRB switch

- create an Ethernet virtual switch with mode “vpls”
- associate the virtual switch with an IP interface

**Note:** The virtual switch cannot be deleted if there are still member UNI AC and NNI PW associating with the VS. The VPLS IRB virtual switch cannot be deleted when it is still referenced by an IP interface.

| Step | Action                                                                                                                                                                                                                                                                                                                  |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Configure an MPLS VPLS IRB virtual switch:</p> <pre>virtual-switch ethernet create vs &lt;String[1..31]&gt; mode vpls reserved-vlan &lt;NUMBER: 1..4094&gt;</pre> <p>where</p> <p>vs &lt;String[1..31] is the virtual switch name.</p> <p>reserved-vlan is the reserved VLAN.</p> <p>&lt;NUMBER:<br/>1..4094&gt;</p> |
| 2    | <p>Add the UNI AC to the VPLS IRB instance:</p> <pre>virtual-switch ethernet add vs &lt;String[1..31]&gt; port &lt;port&gt; vlan &lt;vlan&gt;</pre>                                                                                                                                                                     |
| 3    | <p>Add the NNI PW to the VPLS IRB instance:</p> <pre>virtual-switch ethernet attach mpls-vc vplsirbPw vs &lt;vs&gt;</pre> <p>—end—</p>                                                                                                                                                                                  |

### Example

The following example configures an MPLS VPLS IRB virtual switch, and adds the UNI AC and NNI PW to the VPLS IRB instance:

```
virtual-switch-ethernet create vs vplsirbInst mode vpls reserved-vlan 100
virtual-switch ethernet add vs vplsirbinst port 9 vlan 100
virtual-switch ethernet attach mpls-vc vplsirbPw vs vplsirbInst
```

## Procedure 19-6

### Configuring an MPLS VPLS IP interface associated with the VPLS IRB virtual switch

Configure a virtual IP interface and associate it with the VPLS IRB VS.

**Note 1:** Once the VPLS IRB VS is associated with an IP interface, it cannot be associated with a remote management interface.

**Note 2:** Deleting an IP interface remains unchanged with VPLS IRB.

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Enable VPLS IRB:<br><pre>vpls-irb enable</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 2    | Configure an MPLS VPLS IRB IP interface and associate it with a VPLS IRB virtual switch:<br><pre>interface create ip-interface vs &lt;String[1..31]&gt; ip &lt;ip address&gt; vs &lt;String[1..31]&gt;</pre> <p>where</p> <p>vs &gt;String[1..31] is the virtual switch name of the corresponding VPLS IRB VS instance. This creates a VPLS IRB IP interface. If the VPLS IRB VS is already associated with an IP interface, the configuration is rejected.</p> <p>ip &lt;ip address&gt; is the IP address of the virtual switch.</p> <p style="text-align: center;">—end—</p> |

### Example

The following example configures an MPLS VPLS IRB interface and associates it with a VPLS IRB virtual switch.

```
interface ip-interface vplsirbIf ip 50.50.50.254/24 vs vplsirbinst
```

## Procedure 19-7

### Configuring static ARP over the virtual IP interface

---

The IP interface can be used as a normal IP interface for static route configuration.

---

| Step | Action |
|------|--------|
|------|--------|

---

- 1 Configure the static ARP entry:

```
arp static add destination <IP address> mac <MAC address:
XX:XX:XX:XX:XX:XX> ip-interface <ip-interface>
```

where

destination <IP address> is the destination IP address of the ARP entry.

mac <MAC address: XX:XX:XX:XX:XX:XX> is the MAC address of the MAC entry.

ip-interface <ip-interface> is the name of the IP interface.

—end—

### Example

The following example configures a static ARP over the virtual IP interface associated with the VPLS IRB virtual switch.

```
arp static add destination 50.50.50.2 mac 00:01:02:03:04:05 ip-interface
vplsirbIf
```

## Procedure 19-8

# Configuring OSPF passive interface on the MTU for IP interfaces connecting towards eNBs

Configure OSPF passive interface on the MTU for IP interfaces connecting towards eNBs.

| Step | Action |
|------|--------|
|------|--------|

- |   |                                                                                                                               |
|---|-------------------------------------------------------------------------------------------------------------------------------|
| 1 | <p>Configure OSPF passive interface on the MTU:</p> <pre>ospf interface set ip-interface intf_name passive</pre> <p>—end—</p> |
|---|-------------------------------------------------------------------------------------------------------------------------------|

## Procedure 19-9

# Configuring OSPF area route filtering to separate OSPF route distribution

OSPF area route filtering can be used to separate OSPF route distribution between areas to isolate between IRB and non-IRB. You can also delete an OSPF area.

**Note 1:** Other area networks that fall into the configured prefix will not leak into the area.

**Note 2:** If the area is not created yet, you will not see an error message.

| Step | Action |
|------|--------|
|------|--------|

### To separate OSPF route distribution between areas

- 1 Separate OSPF route distribution between areas to isolate between IRB and non-IRB:

```
ospf area add area-ip (ip_address> prefix-filter
<ip_net_addr/mask> ospf-instance <instance-name>
```

where

area-ip <ip\_address> is the area IP address.

prefix-filter <ip\_addr/mask> is the IP address or mask address that filters out the inter-area prefix.

ospf-instance is the name of the OSPF instance.  
<instance-name>

### To delete an OSPF area

- 2 Delete an OSPF area

```
ospf area delete area_ip <area_ip_addr> prefix-filter
<ip_net_addr/mask> ospf-instance <instance-name>
```

—end—

## Procedure 19-10

### Displaying system reserved ports for VPLS IRB

You can display

- system reserved ports for MPLS VPLS IRB
- system reserved port information

| Step | Action |
|------|--------|
|------|--------|

#### *To display system reserved ports for MPLS VPLS IRB*

- 1 Display system reserved ports for MPLS IRB:  

```
chassis system-rsvd-port show vpls-irb
```

#### *To display system reserved ports information*

- 2 Display system reserved port information:  

```
chassis system-rsvd-port show
```

—end—

### Example

The following example shows output from the chassis system-rsvd-port show vpls-irb command.

```
5160> chassis system-rsvd-port show vpls-irb
```

```
+-----+-----+ SYSTEM RESERVED PORT INFO +-----+
| Feature | Port |
+-----+-----+
| vpls-irb | 24 |
+-----+-----+
```

The following example shows output from the chassis system-rsvd-port show command.

```
5160> chassis system-rsvd-port show
```

```
+-----+-----+ SYSTEM RESERVED PORT INFO +-----+
| Port | Feature |
+-----+-----+
| 24 | vpls-irb |
+-----+-----+
```

## Procedure 19-11

### Displaying the status of MPLS VPLS IRB

Verify that MPLS VPLS IRB is enabled or disabled.

| Step | Action |
|------|--------|
|------|--------|

- |   |                                      |
|---|--------------------------------------|
| 1 | Display the status of MPLS VPLS IRB: |
|---|--------------------------------------|

`vpls-irb show`

—end—

### Example

The following example shows output from the `vpls-irb show` command.

```
> vpls-irb show
+-----+
+ VPLS IRB GLOBAL INFORMATION +
+-----+-----+
Admin Status	Enabled
Oper Status	Enabled
IRB Port	24
+-----+-----+
```

## Procedure 19-12

### Displaying VPLS IRB bandwidth usage statistics

You can monitor bandwidth capacity usage over the loopback port(s). Statistics can be monitored for both the AC and IP interface port on a 3930, 3932 and on the reserved loopback port on a 5160.

| Step | Action |
|------|--------|
|------|--------|

- |   |                                                                                                              |
|---|--------------------------------------------------------------------------------------------------------------|
| 1 | Display VPLS IRB statistics:<br><pre>vpls-irb show statistics</pre> <p style="text-align: center;">—end—</p> |
|---|--------------------------------------------------------------------------------------------------------------|

### Example

The following example shows output from the vpls-irb show statistics command on a 3930 device.

```
3930> vpls-irb show statistics
```

| ----- VPLS IRB AC STATISTICS ----- |             |                  |
|------------------------------------|-------------|------------------|
| Statistic                          | Value       | Rate Mpps & Mbps |
| RxBytes                            | 13243812485 | 30.023           |
| RxPkts                             | 131517901   | 7.010            |
| TxBytes                            | 30008238467 | 77.055           |
| TxPkts                             | 116632759   | 11.032           |

| ----- VPLS IRB IP STATISTICS ----- |             |                  |
|------------------------------------|-------------|------------------|
| Statistic                          | Value       | Rate Mpps & Mbps |
| RxBytes                            | 13738432485 | 32.328           |
| RxPkts                             | 131323901   | 7.555            |
| TxBytes                            | 30332342467 | 123.667          |
| TxPkts                             | 166232759   | 21.001           |



## Procedure 19-13

### Displaying an MPLS VPLS IRB virtual switch

Display the MPLS VPLS IRB virtual switch to display the active VLAN and reserved VLAN information.

| Step | Action |
|------|--------|
|------|--------|

- |   |                                                                                                                                                       |
|---|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | Display an MPLS VPLS IRB virtual switch:<br><pre>virtual-switch ethernet show vs &lt;String[1..31]&gt;</pre> <p style="text-align: center;">—end—</p> |
|---|-------------------------------------------------------------------------------------------------------------------------------------------------------|

### Example

The following example shows the output for the virtual-switch ethernet show command for the vs vplsIrblnst:

| ----- MPLS VIRTUAL SWITCH INFO ----- |             |                      |                      |      |
|--------------------------------------|-------------|----------------------|----------------------|------|
| Parameter                            |             | Value                |                      |      |
| Name                                 | vplsIrbInst |                      |                      |      |
| ID                                   | 4           |                      |                      |      |
| Description                          |             |                      |                      |      |
| Mode                                 | vpls        |                      |                      |      |
| Active VLAN                          | (100)       | VLAN#100             |                      |      |
| Reserved VLAN                        | (100)       | VLAN#100             | (user)               |      |
| Creation                             | Static      |                      |                      |      |
| Transparent Validate                 | off         |                      |                      |      |
| Mac Learning                         | Enabled     |                      |                      |      |
| Subscr Dot1D Policy                  | leave       |                      |                      |      |
| Decap TTL Policy                     | NA          |                      |                      |      |
| Decap CoS Policy                     | leave       |                      |                      |      |
| Decap Fixed Dot1DPri                 | NA          |                      |                      |      |
| Encap CoS Policy                     | fixed       |                      |                      |      |
| Encap Fixed Dot1DPri                 | 2           |                      |                      |      |
| ----- Attachment Circuits -----      |             |                      |                      |      |
| Port                                 | Vlan        | Encap Cos Policy     | Encap Fixed Dot1DPri | Tpid |
| Virtual switch has no members        |             |                      |                      |      |
| ----- Attachment Circuits -----      |             |                      |                      |      |
| Port                                 | Vlan        | Ingress L2 Transform | Egress L2 Transform  |      |
| ----- Attachment Circuits -----      |             |                      |                      |      |
| Tdm Vc                               | Vlan        | Physical Port        |                      |      |
| ----- Virtual Circuits -----         |             |                      |                      |      |
| Name                                 | Decap Label | Encap Label          |                      |      |
| No attached virtual circuits         |             |                      |                      |      |

## Procedure 19-14

### Displaying MPLS VPLS IRB IP interface information

Display an MPLS VPLS IRB IP interface information for

- a specific IP interface
- all IP interfaces

| Step | Action |
|------|--------|
|------|--------|

#### To display information for an MPLS VPLS IRB IP specific IP interface

- 1 Display MPLS VPLS IRB information for a specific IP interface:  

```
interface show ip-interface <interface>
```

#### To display information for all MPLS VPLS IRB IP IP interfaces

- 2 Display MPLS VPLS IRB information for all IRB IP interface:  

```
interface show
```

—end—

### Example

The following example shows the output for the interface show ip-interface command for the Port1Vlan101 interface.

```
> interface show ip-interface Port1Vlan101
```

| IP INTERFACE ENTRY       |                   |
|--------------------------|-------------------|
| Parameter                | Value             |
| Index                    | 2                 |
| Name                     | Port1Vlan101      |
| Admin State              | Enabled           |
| Oper State               | Enabled           |
| RTD Index                | 3                 |
| System IfIndex           | 12                |
| Domain                   | VLAN 101          |
| Priority/PCP             | 7                 |
| MTU                      | 1500              |
| Type                     | Broadcast         |
| IPv4 Forwarding          | on                |
| MAC                      | 9c:7a:03:7d:9f:41 |
| LDP State                | Disabled          |
| Interface RSVP           | Disabled          |
| Tunnel Entities          | Present           |
| L2VPN Association        | Absent            |
| VCCV Profile Association | Absent            |
| IP BFD Session           | Absent            |
| OSPF                     | Detached          |
| ISIS                     | Detached          |
| Static ARP               | Detached          |
| PTP                      | Disabled          |

## 19-24 Virtual private LAN service integrated routing and bridging

| ADMIN INTERFACE ADDRESSES       |                             |        |
|---------------------------------|-----------------------------|--------|
| Parameter                       | Value                       |        |
| IPv4 Addr/Mask                  | 172.16.101.42/24            |        |
| IPv6 Addr/Mask                  | Not configured              |        |
| OPERATIONAL INTERFACE ADDRESSES |                             |        |
| Parameter                       | Value                       | Source |
| IPv4 Addr/Mask                  | 172.16.101.42/24            | Manual |
| IPv6 Addr/Mask                  | fe80::9e7a:3ff:fe7d:9f41/64 | Local  |

The following example shows the output for the interface show command.

> interface show

| INTERFACE MANAGEMENT |                 |                             |
|----------------------|-----------------|-----------------------------|
| Name                 | Domain          | IP Address/Prefix           |
| local                | n/a             | 10.18.2.42/24               |
| local                | n/a             | fe80::9e7a:3ff:fe7d:9f40/64 |
| remote               | VLAN 127        | fe80::9e7a:3ff:fe7d:9f5f/64 |
| lbkIf                | n/a             | 7.7.7.7/32                  |
| Wales-CordorvaA      | VLAN 1800       | 67.14.0.164/24              |
| Wales-CordorvaA      | VLAN 1800       | fe80::9e7a:3ff:fe7d:9f41/64 |
| Wales-CordorvaB      | VLAN 1810       | 67.14.10.164/24             |
| Wales-CordorvaB      | VLAN 1810       | fe80::9e7a:3ff:fe7d:9f41/64 |
| Wales-RenfrewA       | VLAN 1700       | 32.14.0.164/24              |
| Wales-RenfrewA       | VLAN 1700       | fe80::9e7a:3ff:fe7d:9f41/64 |
| Wales-RenfrewB       | VLAN 1710       | 32.14.10.164/24             |
| Wales-RenfrewB       | VLAN 1710       | fe80::9e7a:3ff:fe7d:9f41/64 |
| irb_90               | VS MtuFacing    | 90.90.90.2/24               |
| irb_90               | VS MtuFacing    | fe80::9e7a:3ff:fe7d:9f41/64 |
| irb_41               | VS CoreGwFacing | 41.41.41.2/24               |
| irb_41               | VS CoreGwFacing | fe80::9e7a:3ff:fe7d:9f41/64 |

| TCP/IP/STACK OPERATIONAL STATE |                |
|--------------------------------|----------------|
| Parameter                      | Value          |
| IPv4 Gateway                   | 10.18.2.1      |
| IPv6 Gateway                   | Not configured |
| IPv4 Forwarding                | Off            |
| Default DSCP                   | 0              |
| Mgmt Port Interface            | local          |

| IPV4 STACK STATE            |       |
|-----------------------------|-------|
| Parameter                   | Value |
| ICMP Accept Redirects       | On    |
| ICMP Echo Ignore Broadcasts | Off   |
| ICMP Port Unreachable       | On    |

| IPV6 STACK STATE |         |
|------------------|---------|
| Parameter        | Value   |
| IPv6 Stack       | Enabled |
| Stack Preference | IPv6    |

|                             |     |
|-----------------------------|-----|
| Accept Router Advertisement | On  |
| ICMP Accept Redirects       | On  |
| ICMP Echo Ignore Broadcasts | Off |
| ICMP Port Unreachable       | On  |

---

|                             |       |
|-----------------------------|-------|
| ----- TCP STACK STATE ----- |       |
| Parameter                   | Value |
| TCP Timestamps              | On    |

## Procedure 19-15

### Displaying the static ARP entry over the VPLS IRB IP interface

Display the static ARP entry over the virtual VPLS IRB IP interface.

| Step | Action |
|------|--------|
|------|--------|

- |   |                                                              |
|---|--------------------------------------------------------------|
| 1 | Display the static ARP entry over the VPLS ORB IP interface. |
|---|--------------------------------------------------------------|

```
arp static show
```

—end—

### Example

The following example shows output from the arp static show command.

```
> arp static show
```

| DestinationIp | DMAC              | ifName    | (ifIndex) |
|---------------|-------------------|-----------|-----------|
| 50.50.50.2    | 00:01:02:03:04:05 | vplsIrbIf | (17)      |
| 50.50.50.3    | 00:01:02:03:04:06 | vplsIrbIf | (17)      |

## Procedure 19-16


### Configuration example using integrated routing and bridging and MTUs

This example uses the following nodes:

- MTUs (3930, 3932)
- Aggregate gateway (5160)
- Core gateway (5160)

**Figure 19-1**

**Use case scenario: MTUs (3930, 3932)**



**Note 1:** There is no change to the existing NNI IP interface, MPLS TNL and PW configuration. Therefore, this configuration is not included in this example.

**Note 2:** eNB nodes are likely using static route and specify the MTU-s node as the gateway. There is no need to run OSPF on the virtual IP interface associated with each eNB cluster.

**Note 3:** If it is desired to separate the routing/forwarding of the traffic from the existing IP traffic for management and other purposes, there is a separate area for traffic for management and other purposes.

**Note 4:** The desired implementation for the 3930 and 3932 is to use a hidden internal port for loopback purposes to overcome the hardware limitation, so there is no need to reserve the external faceplate port.

---

| Step  | Action                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <hr/> |                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 1     | <p>Allocate the advanced-l3 resource feature:</p> <pre>resource-manager pool set feature traffic-profiling resource classifier count 256  resource-manager pool set feature traffic-profiling resource counter count 256  resource-manager pool set feature advanced-l3 resource classifier count 256  resource-manager pool set feature advanced-l3 resource counter count 256  resource-manager validate  config save  chassis reboot</pre> |
| 2     | <p>Enable VPLS IRB:</p> <pre>vpls-irb enable</pre>                                                                                                                                                                                                                                                                                                                                                                                            |
| 3     | <p>Configure VPLS IRB VS for each eNB cluster:</p> <pre>vlan create vlan 21 vlan create vlan 22 vlan create vlan 23 vlan add vlan 21 port 1 vlan add vlan 22 port 2 vlan add vlan 23 port 3</pre>                                                                                                                                                                                                                                             |
| 4     | <p>Create the IP interface and associate it with each VLAN:</p> <pre>interface create ip-interface vplsIrbEnbC101 ip 10.1.1.128/24 vlan 21  interface create ip-interface vplsIrbEnbC102 ip 10.1.2.128/24 vlan 22  interface create ip-interface vplsIrbEnbC103 ip 10.1.3.128/24 vlan 23</pre>                                                                                                                                                |
| 5     | <p>Configure VPLS IRB VS for S1 traffic to reach MSC:</p> <pre>virtual-switch ethernet create vs vplsIrbS1 mode vpls reserved-vlan 100</pre>                                                                                                                                                                                                                                                                                                  |

- 6**      **Attach primary and backup PW for the VPLS IRB VS:**  
virtual-switch ethernet attach mpls-vc s1Mtu128PwP vs  
vplsIrbS1  
virtual-switch ethernet attach mpls-vc s1Mtu128PwB vs  
vplsIrbS1
- 7**      **Create the virtual IP interface and associate it with the VPLS IRB VS:**  
interface create ip-interface vplsIrbS1 ip 192.168.3.128/  
24 vs vplsIrbS1
- 8**      **Configure VPLS IRB VS for X2 traffic to reach the aggregate gateway:**  
virtual-switch ethernet create vs vplsIrbX2A mode vpls  
reserved-vlan 50  
virtual-switch ethernet create vs vplsIrbX2B mode vpls  
reserved-vlan 60
- 9**      **Attach the primary and backup PW for the VPLS IRB VS:**  
virtual-switch ethernet attach mpls-vc x2Mtu128PwP vs  
vplsIrbX2A  
virtual-switch ethernet attach mpls-vc x2Mtu128PwB vs  
vplsIrbX2B
- 10**     **Create the virtual switch IP interface and associate it with the VPLS IRB VS:**  
interface create ip interface vplsIrbX2A ip 5.5.5.1/24 vs  
vplsIrbX2A  
interface create ip interface vplsIrbX2B ip 5.5.5.1/24 vs  
vplsIrbX2B
- 11**     **Create OSPF on the virtual IP interface for S1 and X2 traffic:**  
ospf instance create ospf-instance ospfMtu128  
ospf area create area-ip 0.0.0.10 ospf-instance  
ospfMtu128 type normal  
ospf interface attach ip-interface vplsIrbS1 area  
0.0.0.10 ospf-instance ospfMtu128  
ospf interface attach ip-interface vplsIrbX2A area  
0.0.0.10 ospf-instance ospfMtu128  
ospf interface attach ip-interface vplsIrbX2B area  
0.0.0.10 ospf-instance ospfMtu128

—end—



## Procedure 19-17


### Configuration example using MTUs alternative using IP interface

This example uses the following nodes:

- MTUs (3930, 3932)
- Aggregate gateway (5160)
- Core gateway (5160)

**Figure 19-2**

**Use case scenario: MTUs (3930, 3932)**



**Note 1:** There is no change to the existing NNI IP interface, MPLS TNL and PW configuration. Therefore, this configuration is not included in this example.

**Note 2:** eNB nodes are likely using static route and specify the MTU-s node as the gateway. There is no need to run OSPF on the virtual IP interface associated with each eNB cluster.

**Note 3:** There is no desire to separate routing/forwarding of the S1/X2 traffic from the existing IP traffic for management and other purposes. As a result, there is a separate area for S1/X2 traffic for management and other purposes, and for S1/X2 traffic.

**Note 4:** The desired implementation for the 3930 and 3932 is to use a hidden internal port for loopback purposes to overcome the hardware limitation, so there is no need to reserve the external faceplate port.

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Allocate the advanced-l3 resource feature:<br><pre>resource-manager pool set feature traffic-profiling resource classifier count 256 resource-manager pool set feature traffic-profiling resource counter count 256 resource-manager pool set feature advanced-l3 resource classifier count 256 resource-manager pool set feature advanced-l3 resource counter count 256 resource-manager validate config save chassis reboot</pre> |
| 2    | Enable VPLS IRB:<br><pre>vpls-irb enable</pre>                                                                                                                                                                                                                                                                                                                                                                                      |
| 3    | Configure VLAN for each eNB cluster:<br><pre>vlan create vlan 21 vlan create vlan 22 vlan create vlan 23 vlan add vlan 21 port 1 vlan add vlan 22 port 2 vlan add vlan 23 port 3</pre>                                                                                                                                                                                                                                              |

- 4** Create the IP interface and associate it with each VLAN:

```
interface create ip-interface vplsIrbEnbC101 ip
10.1.1.128/24 vlan 21

interface create ip-interface vplsIrbEnbC102 ip
10.1.2.128/24 vlan 22

interface create ip-interface vplsIrbEnbC103 ip
10.1.3.128/24 vlan 23
```
- 5** Configure VPLS IRB VS for S1 traffic to reach MSC:

```
virtual-switch ethernet create vs vplsIrbS1 mode vpls
reserved-vlan 100
```
- 6** Attach primary and backup PW for the VPLS IRB VS:

```
virtual-switch ethernet attach mpls-vc s1Mtu128PwP vs
vplsIrbS1

virtual-switch ethernet attach mpls-vc s1Mtu128PwB vs
vplsIrbS1
```
- 7** Create the virtual IP interface and associate it with the VPLS IRB VS:

```
interface create ip-interface vplsIrbS1 ip 192.168.3.128/
24 vs vplsIrbS1
```
- 8** Configure VPLS IRB VS for X2 traffic to reach the aggregate gateway:

```
virtual-switch ethernet create vs vplsIrbX2A mode vpls
reserved-vlan 50

virtual-switch ethernet create vs vplsIrbX2B mode vpls
reserved-vlan 60
```
- 9** Attach the primary and backup PW for the VPLS IRB VS:

```
virtual-switch ethernet attach mpls-vc x2Mtu128PwP vs
vplsIrbX2A

virtual-switch ethernet attach mpls-vc x2Mtu128PwB vs
vplsIrbX2B
```
- 10** Create the virtual switch IP interface and associate it with the VPLS IRB VS:

```
interface create ip interface vplsIrbX2A ip 5.5.5.1/24 vs
vplsIrbX2A

interface create ip interface vplsIrbX2B ip 6.6.6.1/24 vs
vplsIrbX2B
```

**11 Create OSPF on the virtual IP interface for S1 and X2 traffic:**

```
ospf instance create ospf-instance ospfMtu128
ospf area create area-ip 0.0.0.10 ospf-instance
ospfMtu128 type normal
ospf interface attach ip-interface vplsIrbS1 area
0.0.0.10 ospf-instance ospfMtu128
ospf interface attach ip-interface vplsIrbX2A area
0.0.0.10 ospf-instance ospfMtu128
ospf interface attach ip-interface vplsIrbX2B area
0.0.0.10 ospf-instance ospfMtu128
```

—end—

## Procedure 19-18


### Configuration example using aggregate gateway (5160)

This example uses the following nodes:

- MTUs (3930, 3932)
- Aggregate gateway (5160)
- Core gateway (5160)

**Figure 19-3**

**Use case scenario: aggregate gateway (5160)**



**Note 1:** There is no change to the existing NNI IP interface, MPLS TNL and PW configuration. Therefore, this configuration is not included in this example.

**Note 2:** The 5160 does not have a hidden internal port with sufficient bandwidth available to be used as a reversed loopback port to overcome the hardware limitation. Therefore, it is not necessary to reserve the external face port.

| Step | Action                                                                                                                                                                                                                                                                                        |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Allocate the advanced-l3 resource feature:</p> <pre>resource-manager pool set feature advanced-l3 resource classifier count 512 resource-manager pool set feature advanced-l3 resource counter count 512 resource-manager validate config save chassis reboot</pre>                        |
| 2    | <p>Reserve the external face plate port for loopback:</p> <pre>chassis system-rsvd-port add port 24 feature vpls-irb</pre>                                                                                                                                                                    |
| 3    | <p>Enable VPLS IRB:</p> <pre>vpls-irb enable</pre>                                                                                                                                                                                                                                            |
| 4    | <p>Configure VPLS IRB VS for X2 traffic to reach the core gateway:</p> <pre>virtual-switch ethernet create vs vplsIrbX2CoreGwA mode vpls reserved-vlan 400 virtual-switch ethernet create vs vplsIrbX2CoreGwB mode vpls reserved-vlan 500</pre>                                               |
| 5    | <p>Attach PWs for the VPLS IRB virtual switch:</p> <pre>virtual-switch ethernet attach mpls-vc x2CoreGWPwP vs vplsIrbX2CoreGwA virtual-switch ethernet attach mpls-vc x2CoreGWPwB vs vplsIrbX2CoreGwB</pre>                                                                                   |
| 6    | <p>Create the virtual switch IP interface and associate it with the VPLS IRB VS:</p> <pre>interface create ip-interface vplsIrbX2CoreGwA ip 40.40.40.1/24 vs vplsIrbX2CoreGwA interface create ip-interface vplsIrbX2CoreGwB ip 50.50.50.1/24 vs vplsIrbX2CoreGwB</pre>                       |
| 7    | <p>Configure the VPLS IRB virtual switch for X2 traffic to reach the MTUs:</p> <pre>virtual-switch ethernet create vs vplsIrbX2Mtu mode vpls reserved-vlan 50</pre>                                                                                                                           |
| 8    | <p>Attach the PW for the VPLS IRB virtual switch for each connected MTU:</p> <pre>virtual-switch ethernet attach mpls-vc x2Mtu128Pw vs vplsIrbX2Mtu virtual-switch ethernet attach mpls-vc x2Mtu129Pw vs vplsIrbX2Mtu virtual-switch ethernet attach mpls-vc x2Mtu130Pw vs vplsIrbX2Mtu</pre> |

**9 Create OSPF on the virtual IP interface for S1 and X2 traffic:**

```
ospf instance create ospf-instance ospfAggGw01
ospf area create area-ip 0.0.0.10 ospf-instance
ospfAggGw01 type normal

ospf interface attach ip-interface vplsIrbX2CoreGwA area
0.0.0.10 ospf-instance ospfAggGw01

ospf interface attach ip-interface vplsIrbX2CoreGwB area
0.0.0.10 ospf-instance ospfAggGw01

ospf interface attach ip-interface vplsIrbX2Mtu area
0.0.0.10 ospf-instance ospfAggGw01
```

—end—

## Procedure 19-19


### Configuration example using aggregate gateway (5160)

This example uses the following nodes:

- MTUs (3930, 3932)
- Aggregate gateway (5160)
- Core gateway (5160)

**Figure 19-4**

**Use case scenario: aggregate gateway (5160)**



**Note 1:** There is no change to the existing NNI IP interface, MPLS TNL and PW configuration. Therefore, this configuration is not included in this example.

**Note 2:** The 5160 does not have a hidden internal port with sufficient bandwidth available to be used as a reversed loopback port to overcome the hardware limitation. Therefore, it is not necessary to reserve the external face port.




| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Allocate the advanced-l3 resource feature:</p> <pre>resource-manager pool set feature advanced-l3 resource classifier count 512 resource-manager pool set feature advanced-l3 resource counter count 512 resource-manager validate config save chassis reboot</pre>                                                                                                                                                                |
| 2    | <p>Reserve the external face plate port for loopback:</p> <pre>chassis system-rsvd-port add port 24 feature vpls-irb</pre>                                                                                                                                                                                                                                                                                                            |
| 3    | <p>Enable VPLS IRB:</p> <pre>vpls-irb enable</pre>                                                                                                                                                                                                                                                                                                                                                                                    |
| 4    | <p>Configure VPLS IRB VS for X2 traffic to reach the aggregate gateway:</p> <pre>virtual-switch ethernet create vs vplsIrbX2CoreGw mode vpls reserved-vlan 40</pre>                                                                                                                                                                                                                                                                   |
| 5    | <p>Attach PWs for the VPLS IRB virtual switch for each aggregate gateway:</p> <pre>virtual-switch ethernet attach mpls-vc x2AggreGW01Pw vs vplsIrbX2AggGw virtual-switch ethernet attach mpls-vc x2AggreGWPw vs vplsIrbX2AggGw</pre>                                                                                                                                                                                                  |
| 6    | <p>Create the virtual switch IP interface and associate it with the VPLS IRB VS:</p> <pre>interface create ip interface vplsIrbX2AggGW ip 40.40.40.2/24 vs vplsIrbX2AggGw</pre>                                                                                                                                                                                                                                                       |
| 7    | <p>Create OSPF on the virtual IP interface for X2 traffic:</p> <pre>ospf instance create ospf-instance ospfCoreGw ospf area create area-ip 0.0.0.10 ospf-instance ospfCoreGw type normal ospf interface attach ip-interface vplsIrbX2AggGw area 0.0.0.10 ospf-instance ospfCoreGw type normal ospf interface attach ip-interface vplsIrbX2AggGw area 0.0.0.10 ospf-instance ospfCoreGw</pre> <p style="text-align: center;">—end—</p> |

## Seamless MPLS configuration

Traditional MPLS currently offers several features like Label Distribution Protocol (LDP) extensions for inter-area Label Switched Paths (LSPs), intra- and inter-area Resource Reservation Protocol-Traffic Engineering (RSVP-TE), LDPoverRSVP and pseudowire (PW) switching (multisegment PW) that can be implemented to support end-to-end MPLS services, as depicted in this figure:


**Figure 20-1**  
**Traditional MPLS**



These features do not completely address the flexibility, scalability, resiliency and manageability of end-to-end MPLS networks.

In contrast to traditional MPLS, seamless MPLS is an IETF Working Group initiative that offers a superior alternative for implementing end-to-end MPLS networks. It describes an architecture that can be used to extend MPLS networks to integrate access, aggregation, and core networks into a single MPLS domain, as depicted in the [Seamless MPLS](#) figure.


**Figure 20-2**  
Seamless MPLS



Seamless MPLS is based on existing and well-known protocols, such as Border Gateway Protocol (BGP) and LDP, providing a logical and easy evolution path. It delivers a highly flexible and scalable architecture with the possibility to integrate 100,000s of nodes.

Seamless MPLS supplies end-to-end service-independent transport, thereby removing the need for service-specific configuration in network transport nodes, as depicted in this figure:

**Figure 20-3**  
BGP-LU providing service across multiple AS



Benefits of using seamless MPLS include:


- The seamless MPLS network is one in which all forwarding of packets within the network (from the time a packet enters the network until it leaves the network) is based on MPLS.
- Seamless MPLS allows services to be provisioned wherever they are needed, no matter how the underlying transport is laid out.
- Seamless MPLS extends the core region and integrates metro regions into a single MPLS domain. This single domain makes managing and troubleshooting the transport and services layer more efficient.
- Seamless MPLS makes deploying services faster and more flexible. Existing MPLS networks are typically implemented and operated as separate networks (that is, core, aggregation, and access).
- Seamless MPLS removes domain boundaries and extends the topology into a single MPLS domain, so service can be deployed faster and between any two points (seamless) in the end-to-end MPLS network.
- Seamless MPLS provides the kind of scale today's networks need. Evolving to end-to-end MPLS networks results in networks with a significantly larger number of MPLS nodes.

Seamless MPLS is achieved, for example, for inter-metro Ethernet transport using:

- BGP Labeled Unicast (LU) (BGP-LU) at the access nodes (AN)
- Pseudowire (PW) signaling and attachment circuits

Seamless MPLS technology is utilized as depicted in this figure:

**Figure 20-4**  
**Seamless MPLS technology usage**



The /32 mask in the [Seamless MPLS technology usage](#) figure represents the preferred source IP for the management protocol. The subnet it is part of can be advertised into BGP. For details on configuring the preferred source IP, see “Management interface configuration” in *39XX/51XX Service Delivery, Aggregation and Virtualization Switches Base Configuration*.

Access nodes use BGP-LU. The initiator AN appends the PW header and the BGP-LU header to the outgoing payload. The aggregation node (AGN) swaps the BGP-LU header and appends the RSVP-TE tunnel header to the incoming PW payload.

The following cases are supported:

- Aggregation devices are configured as Interior BGP (iBGP) peers within the aggregation domain.
- Devices interfacing with the Data Communication Network (DCN) gateway are configured as Exterior BGP (eBGP) peers.
- 6500-T participating in the core are configured as eBGP peers.
- eBGP is configured on the aggregation devices that are interfacing with access devices to forward access routes to peers.

**Note:** There is no Open Shortest Path First (OSPF) in the Interior Gateway Protocol (IGP) domain. Only Intermediate System-to-Intermediate System (IS-IS) is supported.

The AGN is also responsible for terminating the RSVP-TE tunnel and swapping the BGP-LU header before sending the frame to the terminating AN. The terminating AN pops the BGP-LU header and the PW header.

The following figures illustrate BGP prefix-independent convergence (PIC):

**Figure 20-5**  
**BGP prefix-independent convergence—A**





**Figure 20-6**  
**BGP prefix-independent convergence—B**



Note the following restrictions:

- BGP for seamless MPLS is supported only on SAOS 6.x 3926m and 3928 platforms.
- All Traffic Engineering (TE) facilities are available only on RSVP-signaled tunnels within an IGP-TE domain (intra-metro).
- BGP LSP segments transiting through IGP take advantage of TE, including:
  - Shared Risk Link Group (SRLG) diversity
  - Link affinity
  - Bandwidth admission control
- End-to-end BGP LSP does not benefit from TE.

## IP prefix lists

An Internet Protocol (IP) prefix list contains one or more ordered entries that are processed sequentially. The evaluation of a prefix against a prefix list ends as soon as a match is found. An IP prefix list specifies a list of networks. When you apply an IP prefix list to a neighbor, the device sends or receives only a route whose destination is in the IP prefix list. The software interprets the prefix lists in order, beginning with the lowest sequence number.



Configuring an IP prefix list and applying it to a neighbor requires CLI commands such as the following:

```
3926> prefix-list create list-name mylist
3926> prefix-list seq add deny list-name mylist 10 10.10.10.2/24
3926> bgp prefix-list set a 100 nbr 10.23.10.10 list-name mylist dir in
```

## Autonomous system path access lists

The autonomous system (AS) path is used to prevent routing loops in BGP. You also can use this routing information to prefer one path to a destination network over another. To set route preferences, configure filters to apply to AS paths when importing and exporting routes. Each autonomous system through which a route passes appends its AS number to the beginning of the AS path.

Ciena uses the AS path access list as the AS path filter in BGP. It is implemented as AS path regular expressions. The content is read from left to right. The left side matches the beginning of the AS path; the right side matches the end of the AS path. Here is sample configuration:

```
3926> as-path access-list create list-name mylist
3926> as-path access-list add list-name mylist action permit regexp ^100
3926> route-map match-as-path add map-name myroutemap access-list-name mylist
seq 10 action permit
```

## Community lists and expanded community lists

A BGP community is a group of routes that share a common property, regardless of their network, autonomous system, or any physical boundaries. In large networks, applying a common routing policy by using prefix lists or access lists requires individual peer statements on each networking device. Using the BGP COMMUNITIES attribute, BGP speakers with common routing policies can implement inbound or outbound route filters based on the community tag, rather than consult long lists of individual permit or deny statements. A COMMUNITIES attribute can contain multiple communities.

A route can belong to multiple communities. The network administrator defines the communities to which a route belongs. By default, all routes belong to the general Internet community.

A BGP community list is used to create groups of communities to use in a match clause of a route map. A community list can be used to control which routes, for example, are accepted, preferred, distributed, or advertised. A community list can also be used to set, append, or modify the communities of a route.



Community lists that can be used are:

- Standard community lists—to specify well-known communities and community numbers.
- Expanded community lists—to filter communities using a regular expression. Regular expressions are used to specify patterns to match COMMUNITIES attributes.

Here is sample configuration:

```
3926> community-list-standard create list-name 1
3926> community-list-standard add list-name 1 action permit value 1:100
3926> route-map match-community add map-name abc seq 10 community-list-name 1
action permit exact-match
```

## **BGP route map**

A BGP route map is a named set of match conditions and parameter settings that the device can use to modify route attributes and to control redistribution of the routes into other protocols. A route map consists of a sequence of instances, the equivalent of rows in a table. The device evaluates a route according to route map instances in ascending numerical order. The route is first compared against instance 1, then against instance 2, and so on. When a match is found, the device stops evaluating the route.

Route maps can contain match clauses and set statements. Each route map contains a permit or deny statement for routes that match the match clauses:

- If the route map contains a permit statement, a route that matches a match statement is permitted; otherwise, the route is denied.
- If the route map contains a deny statement, a route that matches a match statement is denied.
- If a route does not match any match statements in the route map, then the route is denied. This is the default action. To change the default action, configure the last match statement in the last instance of the route map to permit any.
- If there is no match statement, the software considers the route to be a match.
- For route maps that contain address filters, autonomous system (AS) path filters, or community filters, if the action specified by a filter conflicts with the action specified by the route map, the route map action takes precedence over the filter action.

If the route map contains set statements, routes that are permitted by the route map match statements are modified according to the set statements.

For routes that match all of the match statements, the route map set statements can perform one or more of the following modifications to the route attributes:

- Set the community attributes.
- Set the local preference.
- Set the MED (metric).
- Set the IP address of the next-hop device.

When parameters are configured for redistributing routes into BGP4, one of the optional parameters is a route map. If a route map is specified as one of the redistribution parameters, the device matches the route against the match statements in the route map. If a match is found and if the route map contains set statements, the device sets the attributes in the route according to the set statements.

To create a route map, instances of the map can be defined by a sequence number.

Operands for this command include:

- The “permit | deny” options specify the action the device takes if a route matches a match statement:
  - If deny is specified, the device does not advertise or learn the route.
  - If permit is specified, the device applies the match and sets clauses associated with this route map instance.
- The seq-no parameter specifies the instance of the route map being defined. The following illustrates a creation of a route-map instance 10:

```
3926> route-map create map-name test 10
SHELL PARSER FAILURE: '10' - no matching entry found

3926> prefix-list create list-name dcn-list
3926> prefix-list add list-name dcn-list seq 10 action permit ipaddress/mask
2.2.2.0/24
3926> prefix-list add list-name dcn-list seq 4294967295 action deny ipaddress/
mask 0.0.0.0/0
3926> route-map create map-name dcn-map
3926> route-map match-address add map-name dcn-map seq 10 action permit
prefix-list-name dcn-list
```

Additional BGP route map configuration includes setting the accumulated IGP (AIGP) metric. Here is sample configuration:


```
3926> route-map set-metric add map-name myroutemap seq 10 action permit value
100
```

The BGP route flap dampening can also be set. Here is sample configuration:

```
3926> route-map set-dampening add map-name myroutemap seq 10 action permit
```

## BGP Outbound Route Filtering (ORF) capability

The BGP4 Outbound Route Filtering (ORF) capability is used to minimize the number of BGP updates sent between BGP peers.



When the ORF feature is enabled, unwanted routing updates are filtered out, reducing the amount of system resources required for generating and processing routing updates. The ORF feature is enabled through the advertisement of ORF capabilities to peer routers. The locally configured BGP4 inbound prefix filters are sent to the remote peer so that the remote peer applies the filter as an outbound filter for the neighbor.

The ORF feature can be configured with send and receive ORF capabilities. The local peer advertises the ORF capability in send mode, indicating that it accepts a prefix list from a neighbor and applies the prefix list to locally configured ORFs. The local peer exchanges the ORF capability in send mode with a remote peer for a prefix list that is configured as an inbound filter for that peer locally. The remote peer only sends the first update once it receives a ROUTEREFRESH request or BGP ORF with IMMEDIATE from the peer. The local and remote peers exchange updates to maintain the ORF on each router.

If L2-VPN is configured, Blue Planet MCP updates the prefix list and triggers the RouteRefresh request through CLI/Blue Planet MCP.

### ORF usage in the Ciena solution

On the AN that is a local peer, a prefix list is first created with a prefix list name and acceptance of all routes from the remote BGP peer that also happens to be an AGN is denied. The local peer (AN) advertises the ORF capability in

send mode. The remote peer (AGN) receives the ORF capability in receive mode and applies the filter as an outbound policy. The local and remote peers exchange updates to maintain the ORF on each router. Updates are exchanged between peer routers by address family, depending on the ORF prefix list capability that is advertised. This helps reduce the number of system resources required for generating and processing routing updates sent by the BGP peer (AGN). For example:

```
3926> prefix-list create list-name pw-list
!! By default deny all prefixes that will be advertised from AGN
!! Create a rule with maximum seq number so that further rules can be added
with lower sequence number
3926> prefix-list add list-name pw-list seq 4294967295 action deny ipaddress/
mask 0.0.0.0/0
!! Set prefix-list for 'in' direction
3926> bgp prefix-list add as <asn-AN> nbr 200.200.2.154 afi ipv4 safi unicast
listname pw-list dir in
3926> bgp prefix-list add as <asn-AN> nbr 200.200.2.154 afi ipv4 safi
labeledunicast list-name pw-list dir in
!! Set outbound-route-filtering capability for the neighbor
3926> bgp orf add as <asn-AN> nbr 200.200.2.154 afi ipv4 safi unicast orf-type
prefix-list dir tx
3926> bgp orf add as <asn-AN> nbr 200.200.2.154 afi ipv4 safi labeled-unicast
orftype prefix-list dir tx
```

Once the L2-VPN is configured, BluePlanet MCP (or CLI-driven) updates the prefix list to accept route updates for the PW's configured destination IP address. The remote peer (AGN) starts sending these updates to the AN after a route refresh has been configured with the "clear ip bgp" command or after an ORF prefix list with immediate status is processed. For example:

```
!! create pw over bgp tunnel to dest-ip 3.3.3.3
3926> mpls l2-vpn create dynamic-vc pw-1 peer 3.3.3.3 pw-id 100 bgp-tunnel
!! add the pw dest-ip to the prefix-list
3926> prefix-list add list-name pw-list seq 10 action permit ipaddress/mask
3.3.3.3/32
!! soft route reset
3926> bgp clear neighbor nbr 200.200.2.154 in prefix-filter afi ipv4 safi
unicast
3926> bgp clear neighbor nbr 200.200.2.154 in prefix-filter afi ipv4 safi
labeled-unicast
```

To understand how ORF fits into the overall Ciena solution, refer to ["BGP-LU solution \(AN to AN\)" on page 20-13](#).

## ORF generic CLI samples

Here is sample CLI for configuring ORF:

```
3926> prefix-list create list-name mylist
3926> prefix-list seq add deny list-name mylist 10 10.10.10.2/24
3926> bgp create as 100
3926> bgp neighbor add as 100 nbr 100.100.100.1 remote-as 100
3926> bgp address-family add as 100 afi ipv4 safi unicast
3926> bgp next-hop-self enable as 100 nbr 100.100.100.1
3926> bgp prefix-list set as 100 nbr 100.100.100.1 list-name mylist dir in
```

Here is sample CLI for configuring ORF in receive mode:

```
3926> bgp orf enable as 100 nbr 100.100.100.1 orf-type prefix-list dir rx
```

Here is sample CLI for configuring ORF in send mode:

```
3926> bgp orf enable as 100 nbr 100.100.100.1 orf-type prefix-list dir tx
```

Here is sample CLI for configuring ORF in both send and receive modes:

```
3926> bgp orf enable as 100 nbr 100.100.100.1 orf-type prefix-list dir both
```

### BGP fast reroute's role in the BGP PIC feature

BGP Fast Reroute (FRR) provides a best path and a backup/alternate path in BGP. BGP FRR provides a very fast reroute mechanism on the backup BGP next hop to reach a destination when the current best path is not available. BGP FRR precomputes a second best path in BGP and gives it to the datapath as a backup/alternate path, and the datapath programs it into line cards. Therefore, BGP FRR sets up the best path and backup/alternate path. The BGP prefix-independent convergence (PIC) feature provides the ability for SAOS to quickly switch the traffic to the other egress ports if the current next hop or the link to this next hop goes down.

If failure in the peer is detected by IGP, it may take a few seconds to detect the failure. Convergence can occur in subseconds or seconds, depending on whether PIC is enabled on the line cards.

If the failure is with directly connected neighbors and if we use BFD to detect when a neighbor has gone down, the detection happens within a subsecond and the convergence can occur in subseconds or seconds, depending on whether PIC is enabled on the line cards. BGP PIC is supported at the encapsulation and also at the transit nodes.

Here is sample configuration to enable FRR in BGP PIC:

```
3926> bgp create as 10
3926> bgp frr add as 10 afi ipv4 safi labeled-unicast
3926> bgp frr add as 10 afi ipv4 safi unicast
```

### BGP OAM

For BGP OAM, supported functionality includes:

- BGP-LU tunnels support BGP tunnel ping and virtual circuit (VC) ping over BGP-LU tunnels.
- BGP-LU tunnels support BGP trace-route functionality.
- BGP uses multihop Internet Protocol Bidirectional Forwarding Detection (IP-BFD) and single-hop IP-BFD to detect faults and recompute other paths.

Here is sample configuration to enable IP-BFD:

- For multihop iBGP BFD between AGNs, the configuration is:

```
bgp bfd-monitor enable as <as-number> nbr <bgp peer ip> multihop
```

- For single-hop eBGP BFD between AGNs to ANs, the configuration is:

```
bgp bfd-monitor enable as <as-number> nbr <eBGP peer ip>
```

There are global BFD profiles from where BFD sessions can pick up the transmit and receive intervals. These can be configured in the following ways:

- For single-hop BGP BFD sessions:

```
bfd profile set profile IP-BGP-Singlehop-Default transmit-interval 50msec
receive-interval 20msec
```

- For multihop BGP BFD sessions:

```
bfd profile set profile IP-BGP-Multihop-Default transmit-interval 20msec
receive-interval 20msec
```


Other CLI commands that can be used to set 20- and 50-millisecond transmit and receive intervals on BFD profiles are:

```
bfd profile set profile IP-IGP-Default transmit-interval 50msec receive-
interval 20msec
bfd profile create profile new
bfd profile set profile new transmit-interval 20msec receive-interval 50msec
```

## BGP-LU solution (AN to AN)

The [BGP-LU solution \(AN to AN\)](#) figure shows how the BGP tunnel is set up at various nodes when a service is created from AN to AN for a simple four-node topology.

**Figure 20-7**  
**BGP-LU solution (AN to AN)**



The steps to do this are:

- 1 The user configures a static route to AGN.1 through CLI on AN.1.
- 2 The user configures a static route to AN.1 through CLI on AGN.1.
- 3 The user configures a static route to AGN.2 through CLI on AN.2.
- 4 The user configures a static route to AN.2 through CLI on AGN.2.
- 5 The user configures a prefix filter to filter all inbound routes on AN.1.
- 6 The user configures a prefix filter to filter all inbound routes on AN.2.
- 7 AN.1 informs AGN.1 to filter all routes learned on AGN.1. See [“BGP Outbound Route Filtering \(ORF\) capability”](#) on page 20-10 for more information on how this is implemented.
- 8 AN.2 informs AGN.2 to filter all routes learned on AGN.2.
- 9 AN.1 informs AGN.1 to use AN.1’s local label whenever AGN.1 wants to send traffic to AN.1. For example, in the [BGP-LU solution \(AN to AN\)](#) figure, AN.1 advertises AGN.1 to use label 100 whenever AGN.1 needs to communicate with AN.1.
- 10 AN.2 informs AGN.2 to use AN.2’s local label whenever AGN.2 wants to send traffic to AN.2. For example, in the [BGP-LU solution \(AN to AN\)](#) figure, AN.2 advertises AGN.2 to use label 300 whenever AGN.2 needs to communicate with AN.2.



- 11 AGN.1 informs AGN.2 to use AGN.1's local label whenever AGN.2 wants to send traffic to AN.1. For example, in the [BGP-LU solution \(AN to AN\)](#) figure, AGN.1 advertises AGN.2 to use label 101 whenever AGN.2 needs to communicate with AN.1.
- 12 AGN.2 informs AGN.1 to use AGN.2's local label whenever AGN.1 wants to send traffic to AN.2. For example, in the [BGP-LU solution \(AN to AN\)](#) figure, AGN.2 advertises AGN.1 to use label 201 whenever AGN.1 needs to communicate with AN.2.
- 13 AN.1 sets up its local label obtained in step 9 as a Pop label.
- 14 AN.2 sets up its local label obtained in step 10 as a Pop label.
- 15 AGN.1 sets up its local label obtained in step 11 as a Swap label to AN.1 local label. For example, label 101 is swapped to label 100.
- 16 AGN.2 sets up its local label obtained in step 12 as a Swap label to AN.2 local label. For example, label 201 is swapped to label 300.
- 17 VC is configured on AN.1 with its destination IP being the loopback IP of AN.2.
- 18 VC is configured on AN.2 with its destination IP being the loopback IP of AN.1.
- 19 VC configuration on AN.1 is also succeeded by a prefix list update through CLI on AN.1. AN.1 requests AGN.1 to filter all routes except the one belonging to the loopback of AN.2.
- 20 VC configuration on AN.2 is also succeeded by a prefix list update through CLI on AN.2. AN.2 requests AGN.2 to filter all routes except the one belonging to the loopback of AN.1.
- 21 AGN.1 informs AN.1 that to reach AN.2, it needs to use AGN.1's locally generated label. For example, in the [BGP-LU solution \(AN to AN\)](#) figure, AGN.1 advertises AN.1 to use label 401 whenever AN.1 needs to communicate with AN.2.
- 22 AGN.2 informs AN.2 that to reach AN.1, it needs to use AGN.2's locally generated label. For example, in the [BGP-LU solution \(AN to AN\)](#) figure, AGN.2 advertises AN.2 to use label 501 whenever AN.2 needs to communicate with AN.1.
- 23 AGN.1 sets up forward swap label translation based on step 21. For example, label 401 is swapped to 201.
- 24 AGN.2 sets up forward swap label translation based on step 22. For example, label 501 is swapped to 101.


Since there is no guarantee that AGNs are always connected back-to-back, an RSVP-TE tunnel exists between the aggregation nodes. The BGP transit tunnel created at AGN.1 and AGN.2 is a part of the RSVP-TE tunnels between AGN.1 and AGN.2. In other words, AGN.1 advertises the labels to AN.1, only when reachability to AGN.2 is determined through the RSVP-TE tunnel in the forward direction. Ditto for AGN.2, where AGN.2 advertises the labels to AN.2, only when reachability to AGN.1 is determined through the RSVP-TE tunnel in the reverse direction.



## BGP-LU solution (AGN to AN)

The [BGP-LU solution \(AGN to AN\)](#) figure shows how the BGP tunnel is set up at various nodes when a service is created from AGN to AN for a simple four-node topology.

**Figure 20-8**  
**BGP-LU solution (AGN to AN)**



The steps to do this are:

- 1 The user configures a static route to AGN.2 through CLI on AN.2.
- 2 The user configures a static route to AN.2 through CLI on AGN.2.
- 3 The user configures a prefix filter to filter all inbound routes on AN.1.
- 4 The user configures a prefix filter to filter all inbound routes on AN.2.
- 5 AN.2 informs AGN.2 to filter all routes learned on AGN.2. See [“BGP Outbound Route Filtering \(ORF\) capability”](#) on page 20-10 for more information on how this is implemented.
- 6 AN.2 informs AGN.2 to use AN.2’s local label whenever AGN.2 wants to send traffic to AN.2. For example, in the [BGP-LU solution \(AGN to AN\)](#) figure, AN.2 advertises AGN.2 to use label 300 whenever AGN.2 needs to communicate with AN.2.
- 7 AGN.2 informs AGN.1 to use AGN.2’s local label whenever AGN.1 wants to send traffic to AN.2. For example, in the [BGP-LU solution \(AGN to AN\)](#) figure, AGN.2 advertises AGN.1 to use label 201 whenever AGN.1 needs to communicate with AN.2.
- 8 AN.2 sets up its local label obtained in step 10 as a Pop label.


- 9 AGN.1 sets up its local label advertised to AGN.2 as a Pop label. For example, label 101 is popped.
- 10 AGN.2 sets up its local label obtained in step 12 as a Swap label to AN.2 local label. For example, label 201 is swapped to label 300.
- 11 VC is configured on AGN.1 with its destination IP being the loopback IP of AN.2.
- 12 VC is configured on AN.2 with its destination IP being the loopback IP of AGN.1.
- 13 AGN.1 identifies that to reach AN.2, it needs to use AGN.2's advertised label. For example, in the [BGP-LU solution \(AGN to AN\)](#) figure, AGN.2 advertises AGN.1 to use label 201 whenever AGN.1 needs to communicate with AN.2.
- 14 AGN.2 informs AN.2 that to reach AGN.1, it needs to use AGN.2's locally generated label. For example in the [BGP-LU solution \(AGN to AN\)](#) figure, AGN.2 advertises AN.2 to use label 501 whenever AN.2 needs to communicate with AN.1.
- 15 AGN.1 sets up forward encapsulation label translation based on step 21. For example, label 201 is configured as a Push label.
- 16 AGN.2 sets up forward swap label translation based on step 22. For example, label 501 is swapped to 101.

Since there is no guarantee that aggregation nodes are always connected back-to-back, an RSVP-TE tunnel exists between the aggregation nodes. The BGP ingress tunnel created at AGN.1 and transit tunnel created at AGN.2 are a part of the RSVP-TE tunnels between AGN.1 and AGN.2. In other words, AGN.1 sets up the encapsulation BGP tunnel only when reachability to AGN.2 is determined through the RSVP-TE tunnel in the forward direction. Ditto for AGN.2, where AGN.2 advertises the labels to AN.2, only when reachability to AGN.1 is determined through the RSVP-TE tunnel in the reverse direction.

## BGP-LU use cases

Expanding on the configuration example discussed in "[BGP-LU solution \(AN to AN\)](#)" on page 20-13 and "[BGP-LU solution \(AGN to AN\)](#)" on page 20-16, the [BGP-LU use cases](#) figure shows the possible use cases where BGP-LU tunnels are set up.

**Figure 20-9**  
**BGP-LU use cases**



Use case 1 shows the scenario where a service is created between the aggregation node AGN5 and the access node AN. Here the BGP tunnel is set up between the aggregation node and the access node. The RSVP-TE tunnel is set up between the aggregation nodes AGN5 and AGN6.

The VC at AGN5 exists over a BGP tunnel, which in turn uses an RSVP-TE tunnel.

The VC at AN exists over a BGP tunnel.

Use case 2 shows the scenario where a service is created between the access node AN1 and the access node AN2. Here the BGP tunnel is set up between the access node AN1 and the access node AN2. The RSVP-TE tunnel is set up between the aggregation nodes AGN1 and AGN2.

The BGP transit tunnels at AGN1 and AGN2 use an RSVP-TE tunnel for forward and reverse directions.

The VC at AN1 and AN2 exist over a BGP tunnel.

Use case 3 shows the scenario where a service is created between two aggregation nodes over a dynamic co-routed tunnel.

## IP over BGP LU features


Seamless MPLS provides highly scalable end-to-end MPLS transport architecture. It can carry various user data, such as VC and Layer 3 IP traffic.

This section shows how the IP management network can be integrated into seamless MPLS transport.

### Management and control traffic

The following figures provide management and control traffic examples for IP over BGP-LU.

**Figure 20-10**  
Management traffic (IP forwarding) AN to AGN (Tx direction)




AN1 builds an eBGP session with a near-end AGN2. AN1 advertises its management and other interfaces as eBGP IPv4 routes into near-end AGN2. AGN2 redistributes these eBGP IPv4 routes as LU routes in the network. All remote AGN nodes can reach AN1 through a BGP-LU route.

AN1 transmits IP management traffic to the remote DCN:


- 1 AN1 sends regular IPv4 packets to AGN2.
- 2 AGN2 forwards packets using a LU tunnel. The top label is one of the RSVP IGP tunnel labels. The bottom label is a BGP-LU label.
- 3 When packets received at the remote LU egress the AGN node, the MPLS is terminated. The IPv4 payload is forwarded to DCN using the IPv4 path.

**Figure 20-11**  
**Management traffic (IP forwarding) AN to AGN (Rx direction)**



When AGN2 receives management IOverMPLS LU tunneled traffic destined for AN3, the MPLS label is terminated. Native IP traffic is forwarded using the IPv4 forward entry path.

**Figure 20-12**  
**Control traffic (IP forwarding) AN to AGN (Tx direction)**




The AN1 MPLS loopback interface route is used as a Targeted LDP (T-LDP) session peering address. It needs to be advertised to AGNs as a BGP LU route.

**Note:** This LU route carries VC data from end-to-end.

Since the AN node does not support IOverMPLS encapsulation, the T-LDP control packets are IPv4 path-forwarded to AGN2. AGN2 then populates the route as the LU route.

**Figure 20-13**  
**Control traffic (IP forwarding) AN to AGN (Rx direction)**



Since AN1 advertises the IP-MPLS loopback interface route as an LU route, AN1 receives T-LDP control packets as MPLS labeled packets.

Therefore, AN1 needs to terminate IOverMPLS label packets. It removes the LU label and forwards T-LDP control packets to the LDP protocol.

### AGN IP active routes selection

A system can have both IPv4 and MPLS BGP-LU routes reaching toward the same destination.

The active routes selection algorithm is based on the route type administrative distance value.

By default, the per route type administrative distance values sorted order (from small to large) is:

- Connected local routes

- IP static routes
- BGP-LU routes
- OSPF routes
- IS-IS routes

A route with the smallest administrative distance value is selected as an active forwarding route.

**Note:** Selected active routes in an AGN are used to forward both self-generated and transit IP traffic.

### Various IP traffic forwarding paths in an AGN

The traffic forwarding paths for various IP packets are:

- OSPF goes through the connected IPv4 path.
- IS-IS goes through the connected IPv4 path.
- RSVP goes through the connected IPv4 path.
- T-LDP goes through the selected active route path.
- BGP Multihop BFD (M-BFD) goes through the selected active route path.
- BGP goes through the selected active route path.
- All other IP unicast packets go through the selected active route path.

**Note:** Since SAOS 6.x switches can be deployed as AN nodes, SAOS 6.x initially supports the BGP-LU label decapsulation feature. Ciena's 6500 platform supports IOverMPLS encapsulation/decapsulation features.

### IP High Availability (HA)

BGP-LU routes HA is supported. All BGP-LU routes in active Control Timing and Switch Module (CTX) are check-pointed into standby.


### IP over MPLS on AN node

This section provides an architecture that supports handling of control traffic over the BGP tunnel on the AN node. There is need of IOverMPLS receipt support for LDP/IP traffic destined to the control loopback IP that arrives on the BGP tunnel.

The [Control and management traffic data path from AN to core/EMS and vice versa](#) figure shows the complete data path followed by control and management traffic from the AN to core/EMS and vice versa.



**Figure 20-14**  
Control and management traffic data path from AN to core/EMS and vice versa



### Supported activities


- 1 Receiving IOverMPLS with IP-DA of control loopback of the AN.
- 2 IP forwarding of IOverMPLS frames received on an AN with IP-DA of some end-user connected through the AN.
- 3 IOverMPLS and MPLS-VC data forwarding over the same BGP tunnel and hence, should co-exist in the HW.
- 4 IP path if BGP tunnel is BOS. If not BOS, it follows the data flow with VPN driven from MPLS-VC label.
- 5 IOverMPLS packet reception from the network and forward these to the kernel on the IP Interface mapped to the transport VLAN.
- 6 IOverMPLS Tx encapsulation on the AN node.



## Frame receipt on AN


Following is the use case for receipt of IPoverMPLS frames from the peer end.

### Figure 20-15 Control traffic (IP forwarding) AGN to AN



Whenever such frames are received, the hardware abstraction layer (HAL) extracts the BGP label and sends the IP frame to the kernel using transport VLAN. From the kernel, the normal IP path is followed, either forwarding the IP frame to the end-user or if it is destined for AN, the packet is consumed by the required application.

**Figure 20-16**  
**IP frame forwarding**




IPOverMPLS encapsulation (Tx) is not planned for this release, so all Tx from the AN follow the IPOverVLAN path.

The AN must have IPOverVLAN connectivity towards its peers.

Following is the use case for the receipt of management frames from the peer AGN. The AN receives IPOverVLAN frames on the control/management VLAN from the uplink interface.

**Figure 20-17**  
**Management traffic (IP forwarding) AGN to AN**



## Pseudowire commands

Existing pseudowire commands are enhanced to support seamless MPLS.

To use BGP label switched paths (LSP), a knob is added to the existing “pw” configuration command as follows:

```
mpls l2-vpn create static-vc pw-id <pw_id> peer <peer_id>
in-label <in_label> out-label <out_label> {bgp-lsp | ldp-lsp |
te-tunnel... | tp-tunnel...} [...]
```

The following “pw show” commands can also be used:

```
mpls bgp-lsp show [<ingress | transit | egress>]
```

```
mpls tunnel show [bgp-lsp]
```

## Label ranges

For seamless MPLS configuration, valid label ranges are:

- MPLS label range: 1 to 495903
- TMD VC label range: 495904 to 499999
- BGP label range: 500000 and above

## Procedure

A seamless MPLS configuration procedure is:

- [“Sample BGP configuration for a three-node topology with PIC” on page 20-28.](#)




## Procedure 20-1

### Sample BGP configuration for a three-node topology with PIC

This figure illustrates a three-node topology with prefix-independent convergence (PIC):

**Figure 20-18**  
A three-node topology with PIC



The traffic flows from the IXIA-1 on the right and is received on the IXIA-2 on the left. This configuration only creates a unidirectional BGP-LU tunnel from B to E. No IS-IS is configured.

### Example

Here is the sample BGP configuration for a three-node topology with PIC:

#### Configuring the device at E

```

bgp create as 100
bgp router-id set as 100 ip 3.3.3.3
bgp neighbor add as 100 nbr 200.200.200.2 remote-as 100
bgp address-family add as 100 afi ipv4 safi unicast
bgp address-family add as 100 afi ipv4 safi labeled-unicast
bgp neighbor set as 100 nbr 200.200.200.2 activate afi ipv4 safi unicast
bgp neighbor set as 100 nbr 200.200.200.2 activate afi ipv4 safi labeled-unicast
bgp allocate-label set as 100 all
bgp redistribute add as 100 type connected
bgp neighbor add as 100 nbr 40.40.40.1 remote-as 100
bgp neighbor set as 100 nbr 40.40.40.1 activate afi ipv4 safi unicast
bgp neighbor set as 100 nbr 40.40.40.1 activate afi ipv4 safi labeled-unicast
bgp neighbor add as 100 nbr 50.50.50.1 remote-as 100
bgp neighbor set as 100 nbr 50.50.50.1 activate afi ipv4 safi unicast
bgp neighbor set as 100 nbr 50.50.50.1 activate afi ipv4 safi labeled-unicast
mpls l2-vpn create static-vc vc1 pw-id 1 peer 1.1.1.1 in-label 7000 out-label 7000 bgp-
tunnel
virtual-switch ethernet create vs vs1 mode vpls
virtual-switch ethernet attach mpls-vc vc1 vs vs1

```

## Configuring the device at A

```

bgp create as 100
bgp router-id set as 100 ip 2.2.2.2
bgp neighbor add as 100 nbr 200.200.200.1 remote-as 100
bgp neighbor add as 100 nbr 100.100.100.1 remote-as 100
bgp address-family add as 100 afi ipv4 safi unicast
bgp address-family add as 100 afi ipv4 safi labeled-unicast
bgp neighbor set as 100 nbr 200.200.200.1 activate afi ipv4 safi unicast
bgp neighbor set as 100 nbr 200.200.200.1 activate afi ipv4 safi labeled-unicast
bgp neighbor set as 100 nbr 100.100.100.1 activate afi ipv4 safi unicast
bgp neighbor set as 100 nbr 100.100.100.1 activate afi ipv4 safi labeled-unicast
bgp allocate-label-all set as 100 all
bgp next-hop-self add as 100 nbr 100.100.100.1 afi ipv4 safi unicast
bgp next-hop-self add as 100 nbr 100.100.100.1 afi ipv4 safi labeled-unicast
bgp route-reflector-client add as 100 nbr 100.100.100.1 afi ipv4 safi unicast
bgp route-reflector-client add as 100 nbr 100.100.100.1 afi ipv4 safi labeled-unicast
bgp bfd-monitor enable as 100 nbr 200.200.200.1
bgp frr add as 100 afi pv4 safi unicast
bgp frr add as 100 afi ipv4 safi labeled-unicast
bgp neighbor add as 100 nbr 50.50.50.2 remote-as 100
bgp neighbor set as 100 nbr 50.50.50.2 activate afi ipv4 safi unicast
bgp neighbor set as 100 nbr 50.50.50.2 activate afi ipv4 safi labeled-unicast

```

## Configuring the device at B

```

bgp create as 100
bgp neighbor add as 100 nbr 100.100.100.2 remote-as 100
bgp address-family add as 100 afi ipv4 safi unicast
bgp address-family add as 100 afi ipv4 safi labeled-unicast
bgp neighbor set as 100 nbr 100.100.100.2 activate afi ipv4 safi unicast
bgp neighbor set as 100 nbr 100.100.100.2 activate afi ipv4 safi labeled-unicast
bgp allocate-label-all set as 100 all
bgp bfd-monitor enable as 100 nbr 100.100.100.2
bgp frr add as 100 afi pv4 safi unicast
bgp frr add as 100 afi ipv4 safi labeled-unicast
bgp neighbor add as 100 nbr 40.40.40.2 remote-as 100
bgp neighbor set as 100 nbr 40.40.40.2 activate afi ipv4 safi unicast
bgp neighbor set as 100 nbr 40.40.40.2 activate afi ipv4 safi labeled-unicast
mpls l2-vpn create static-vc vc1 pw-id 1 peer 3.3.3.3 in-label 7000 out-label 7000 bgp-
tunnel
virtual-switch ethernet create vs vs1 mode vpls
virtual-switch ethernet attach mpls-vc vc1 vs vs1
virtual-switch ethernet add vs vs1 port 3 vlan 40

```

—end—



---

## Flow loop detection for MPLS services

---


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.



**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.

## 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**



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:** Support for flow loop detection is extended for UNI ports configured as attachment circuits of an MPLS service.

### Loop detection through MAC motion support for the MPLS network

In the MPLS network, SMAC learning at the UNI side is done with respect to MPLS AC ports and at the NNI side, learning is done with respect to MPLS VC (PW - Pseudo Wire), which is different from Ethernet services. Here MAC flapping can occur between two or more ACs and/or PWs. Therefore, loop termination action is performed on the MPLS AC port or MPLS VC. The support for loop detection is limited to the UNI side of the MPLS network. The configuration for MPLS AC ports on the UNI side is on a “per port basis”, the same as Ethernet services.

## 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 notify-only, only notifications are sent to the user. Notifications are sent even if the configured actions are Deny MAC movement or Port shutdown.

## 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.

The loop termination/prevention action applied on a port is also applicable on UNI ports configured as attachment circuits in MPLS services.

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.


## Configuration examples for MPLS services

Some examples of MPLS topology and configuration are provided here.

## Loop detection through MAC motion on an MPLS network—Example 1

Figure 21-4

Example 1 topology




### Example 1 configuration

```
flow loop-detection enable
flow learning enable vs vs1
flow loop-detection enable port 1
```

## Loop detection through MAC motion on an MPLS network—Example 2


Figure 21-5

Example 2 topology



### Example 2 configuration


```
flow loop-detection enable
flow learning enable vs vs1
flow loop-detection enable port 1
```

**Loop detection through MAC motion on an MPLS network—Example 3****Figure 21-6****Example 3 topology****Example 3 configuration**

```

flow loop-detection enable
flow learning enable vs vs1
flow loop-detection enable port 1
flow loop-detection enable port 2
flow loop-detection enable port 3

```

**Loop detection through MAC motion on an MPLS network—Example 4****Figure 21-7****Example 4 topology**

### Example 4 configuration

```


flow loop-detection enable
flow learning enable vs vs1
flow learning enable vs vs2
flow loop-detection enable port 1

```

## Loop detection through MAC motion on an MPLS network—Example 5

Figure 21-8

Example 5 topology



### Example 5 configuration

```

flow loop-detection enable
flow learning enable vs vs1
flow learning enable vs vs2
flow loop-detection enable port 1
flow loop-detection enable port 2

```

## Procedures

Procedures for flow loop detection are:

- “Enabling and disabling flow loop detection” on page 21-9
- “Enabling MAC learning at a virtual switch” on page 21-10
- “Configuring flow loop detection” on page 21-11
- “Resetting flow loop detection to default values” on page 21-13
- “Clearing loop detection statistics on the specified port(s)” on page 21-14
- “Displaying flow learning information” on page 21-15
- “Displaying flow loop detection information” on page 21-16

## Procedure 21-1

### Enabling and disabling flow loop detection

You must first enable flow loop detection globally, before enabling it on MPLS AC 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 enabled on all MPLS AC 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 MPLS AC UNI ports*

- 2 Enable flow loop detection on MPLS AC UNI ports:  

```
flow loop-detection enable port [<port>]
```

where

port <port> is the MPLS AC UNI port you are enabling flow loop detection functionality on.

#### *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
```



## Procedure 21-2

### Enabling MAC learning at a virtual switch

---

Make sure that MAC learning at a virtual switch is enabled. By default, MAC learning is enabled at the virtual switch. If it is not enabled, enable it.

---

| Step | Action |
|------|--------|
|------|--------|

---

- |   |                                                             |
|---|-------------------------------------------------------------|
| 1 | Make sure that MAC learning is enabled on a virtual switch: |
|---|-------------------------------------------------------------|

```
flow learning enable vs <vs-name>
```

where

vs <vs-name> is the VS MAC learning is enabled on.

—end—

### Example

This example enables MAC learning on vs vs1.

```
flow learning enable vs vs1
```

## Procedure 21-3

### Configuring flow loop detection

You can configure flow loop detection:

- Optionally configure loop prevention action on MPLS AC 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 MPLS AC UNI port(s)

- 1 Optionally configure loop prevention action on MPLS AC UNI port(s):  

```
flow loop-detection set port <port> action <block-mac-
movement | notify-only | port-shut>
```

where

port <port> is the MPLS AC 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.

notify-only sends notifications (SNMP traps/syslog event) only.

port-shut makes the port operationally down.

#### 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.

**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
```

## 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—

## 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):<br>flow loop-detection clear port <port> statistics<br>—end— |
|---|----------------------------------------------------------------------------------------------------------------------|

## 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 ----- |           |      |  |                      |
|-------------------------------|-----------|------|--|----------------------|
| ID                            | Name      | VLAN |  | Learning Status      |
|                               |           |      |  | 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 ----- |        |          |          |          |
|-----------------------------|--------|----------|----------|----------|
| VS                          | Active | Learning | Status   |          |
|                             | VLAN   | Admin    |          | Oper     |
| vsE1                        | 101    | Disabled | Disabled | Disabled |
| vsE2                        | 102    | Enabled  | Enabled  | Enabled  |

## 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 |
+-----+-----+
```

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 |





---

## Loop-free alternate fast reroute configuration

---


Loop-free alternate (LFA) and remote loop-free alternate (R-LFA) are fast reroute (FRR) functionalities used to achieve resiliency and fast protection switching in Internet Protocol (IP) and MPLS Label Distribution Protocol (LDP) networks without creating forwarding loops.

In an IP/MPLS network, Interior Gateway Protocol (IGP) is used to advertise the IP information between routers and each router calculates the best path to possible destinations. LDP is used to distribute MPLS labels to create label switched paths (LSPs) between routers. LDP depends on IGP to understand the network topology and create LSPs.

One of the problems in IP/MPLS network is that when a failure occurs, then both IP and MPLS traffic going through the failed resource suffers the loss until the IGP converges and new paths are calculated. The IGP convergence after a failure can take a long time and can affect the applications that are running on the underlying IP/MPLS network.

This figure depicts a sample failure in an IP/MPLS network:

**Figure 22-1**  
**Sample failure in an IP/MPLS network**



In this figure, node S has a route to destination D through next hop N. All the IP and LDP LSPs follow the path from S to D through next hop N. When the link between S and N fails, or node N itself fails, all traffic from S to D drops, until the IGP converges back and S calculates a new path to reach D. This IGP convergence and new path calculation can take a very long time, depending on topology and IGP timer configurations.

IP/LDP LFA is used to overcome such bottlenecks in an IP/MPLS network. IP/LDP LFA is also called IP/LDP fast reroute (IP FRR or LDP FRR). In this chapter, the concepts applied to IP LFA/R-LFA are also applied to LDP. The term IP LFA/R-LFA also includes LDP LFA/R-LFA.

## Loop-free alternate


IP FRR or IP LFA provides a solution to avoid the long traffic loss due to failure and provides better resiliency to the IP/MPLS network. The LFA solution achieves the goal of faster traffic restoration by pre-calculating an alternate next hop for a specified destination. In case the primary next hop link or primary next hop node fails, the alternate next hop is used to reroute the packets to destination. The alternate next hop is also referred to sometimes as the backup next hop. LFA specifies no change or addition or update in any IGP or LDP protocol. The LFA process is a local node calculation and local decision. It is possible that in a network, a few nodes use the LFA process and a few do not.

LFA can also be used to provide:

- Link protection—to prevent failure of the primary link to next hop router
- Node protection—to prevent both link and node failure of the next hop router

The alternate next hop to be selected is based on various rules. The key rule is that it should be loop-free. For example, in the [Loop caused by alternate](#) figure, if router S selects router A as the alternate next hop, then it can create a loop. When S redirects the traffic to router A, router A forwards the traffic back to node S, causing a loop.

**Figure 22-2**  
**Loop caused by alternate**




The loop-free alternate or LFA selects the alternate next hop that does not cause a loop.

If the example topology shown in the [Loop caused by alternate](#) figure had a slight difference, then node A would have been a LFA next hop. The [LFA providing link protection](#) figure shows a case where node A provides a LFA and at the time of failure, a traffic loop does not occur.

### Figure 22-3

#### LFA providing link protection




The solution shown in the [LFA providing link protection](#) figure works fine for link protection, but if the primary next hop node N fails, then the LFA node B cannot provide a suitable backup.

The node protection LFA selects the alternate next hop such that it provides against node failure, too.

This figure depicts a LFA that is providing protection against node failure:

### Figure 22-4 LFA providing node protection



To solve the example problem, node S selects node B as its alternate next hop, since traffic going for D to B does not cause a loop.


The traffic shown in the [LFA providing link protection](#) figure represents both IP routed and MPLS traffic (LDP LSP). In LDP FRR, the LFA provides the backup label and backup next hop, so that MPLS traffic from S can be routed to the alternate next hop with the label distributed by the alternate next hop for FEC D.

The LFA functionality depends on topology. In some topologies, loop-free alternates are not possible, backup is never provided, and overall coverage of LFA is very low. To provide more coverage for LFA, the concept of remote loop-free alternate (R-LFA) is introduced, although it is still possible that in some topologies, neither LFA nor R-LFA are able to provide protection.

## Remote loop-free alternate

LFA provides a loop-free alternate and protects against link or node failure of the primary next hop. Although in some topologies as shown in this figure, LFA protection is not possible.

**Figure 22-5**  
**LFA protection not possible**




In the topology shown in the [LFA protection not possible](#) figure, node S has a route to destination D via link S-N, and N is the primary next hop for this route. In this topology the other next hop of S is node A. Node A cannot be selected as an alternate next hop, as it does not provide the loop-free condition.

Remote loop-free alternate (R-LFA) provides a solution for topologies where LFA is not possible. The idea of R-LFA is to transport or tunnel the packets to a remote node in the network, so that when packets reach that remote node, they get forwarded to their destinations without a loop.

The R-LFA solution uses some calculations to select the remote node. Once the remote node is selected, the LDP LSP created between S and the remote node is used to tunnel the packets to the remote node. This way, when a fault occurs on the primary path, the S node tunnels the packets to the remote node using the LDP LSP to reach the remote node. Once packets reach the remote node, the existing routing/forwarding entries on the remote node forward the traffic to its destination.

**Figure 22-6**  
**Remote LFA**



The R-LFA uses a few standard terms in its calculation, as explained here:

**Note:** All calculations for LFA/R-LFA are done on the source node S.

- P-Space—This represents the set of the nodes reachable from the source node without traversing the S-N link.
- Extended P-Space—This represents the P-Space of the neighbor nodes of the source node. In the [Remote LFA](#) figure, this includes the P-Space of S and the P-Space of A. This covers nodes A, B, C, and E.

- Q-Space—In relation to a protected link, a router's Q-Space is the set of routers in which that particular router can be reached with no paths traversing that protected link. In the [Remote LFA](#) figure, nodes D, E, and C are part of this set.
- PQ node—This is the set of nodes that are both in Extended P-Space and Q-Space. In the [Remote LFA](#) figure, nodes C and E are part of this set.

In remote LFA, the remote node is selected from the PQ node set. Generally, the PQ node with a lower metric from the source node is chosen as the final remote node. In the [Remote LFA](#) figure, node C is selected as the remote node.

The remote LFA functionality then uses the LDP LSP between S and C as the virtual interface for tunneling the packets from source node S to remote node C in case the primary next hop fails.


The source node also creates a T-LDP session with remote node C, to get the labels needed by C to reach to destination D. This label is used for forwarding the MPLS traffic to node C in case of failure.

The MPLS traffic in a normal situation follows the S-N-D path. When the primary path fails, the MPLS packets are added a backup label provided by C for D and then tunneled inside the tunnel to C. When these MPLS packets are popped at node C, they forward the MPLS frame on the basis of the inner label for D.

This figure depicts the flow for R-LFA in case of primary link failure:



**Figure 22-7**  
**R-LFA flow for IP and MPLS traffic**



As shown in the [R-LFA flow for IP and MPLS traffic](#) figure, the backup next hop is decided as R-LFA with remote node C.

When a fault occurs, the IP and MPLS traffic is tunneled inside the LDP LSP from S to C. The C node then forwards the traffic according to its own forwarding table.

## Reversion


The LFA feature does not support any manual reversion or wait-to-revert timer. As soon as IGP converges and the new primary next hops are selected and programmed, the data shifts to the new primary path.

## Micro-loop prevention

In a network, transitory micro-loops can occur. These cannot be avoided fully. One such transitory micro-loop is possible when the IGP re-converges after a failure. When after a failure, a node converges its IGP before its peer node, it may happen that the nodes in the new converged path are not yet updated, causing a micro-loop until the nodes get their IGP converged.

This figure shows such a case.

**Figure 22-8**  
**Micro-loop situation after convergence**



The support of micro-loop prevention in such cases is done by delaying the IGP convergence on the local link down using a defined timer. The timer is configurable. It must be set to value that allows the whole network to complete the IGP updates. For example, if the network has 15 nodes and a single failure update takes 10 sec. in total to be propagated to all nodes in the network, then the timer value should be kept greater than or equal to 10 sec. Thus, at the node where a failure has been detected, the updated routes will be pushed to the routing table only after 10 sec., at which time all other nodes will have updated their routing tables, and this will avoid any looping in the network.

In the [Micro-loop situation after convergence](#) figure, on the link S-N failure, node S delays its IGP convergence by a timer called “ULOOP\_DELAY\_DOWN TIMER”. This way, node S delays its IGP updates until the other nodes have already updated, thus preventing the micro-loop. Note this functionality only avoids micro-loops in some conditions and does not guarantee that in all situations micro-loops do not occur.

## Reserved IP address for loop-free alternate

The LFA functionality on SAOS 6.x devices requires two internal IP addresses to work. These internal IP addresses must be from a private IP space and with the /31 subnet. These IP addresses must not be used anywhere in the network for any other purpose. Note the same addresses can be used on all nodes in the network for LFA internal purposes. For example, in some networks, if the subnet 192.168.100.2/31 is not used anywhere in the network, then this 192.168.100.2/31 can be set as the LFA reserved IP address on all nodes of the network where LFA is to be used.

If the reserved IP address for LFA is not set, then the system works with the default reserved IP address (172.16.233.2/31), which may conflict if used somewhere else in the network and can further impact traffic.

## Configuration and operations

SAOS 6.x support of LFA requires some configuration on devices and has some variations as listed here:

- SAOS 6.x devices that do not support hardware-based switchovers and rely on software may not perform as fast at LFA FRR switchovers as those that support hardware-based switchovers.
- IP LFA or IP FRR is software-based. The backup part of the routes is kept in the software, and when failure occurs, is pushed in the hardware to shift traffic to the backup path. Switchover performance for IP FRR is less than that of MPLS FRR.
- When enabled, LFA is enabled by default for all interfaces.
- The LFA is only supported on point-to-point IP interfaces. So before enabling LFA, the operator should make the IP interfaces point-to-point. Ciena recommends to keep LFA enabled on all interfaces of a node. If some interfaces are not point-to-point or LFA is disabled on them, then LFA functionality may not give the desired protection.
- In SAOS during convergence, it is possible that some minimal traffic drop occurs due to hardware limitations and other network and protocol convergence.
- There can be a traffic drop after convergence due to micro-loop. It depends on other nodes and their IGP updates.
- The LFA calculation delay timer should be set higher than the micro-loop prevention timer. Ciena recommends to keep the LFA calculation delay timer 2 to 3 seconds higher than the micro-loop prevention timer.
- The IP traffic that is routed through the kernel may not be able to meet the same performance during switchover as that of hardware-switched traffic. IP traffic coming from the kernel may resume after IGP convergence, in some cases and in some platforms.
- LFA operates independently in each contiguous IGP area. In IS-IS LFA for a prefix, if primary and alternate are within the same IS-IS level, then LFA can be provided for that prefix.
- Other protocols and features listed here are not supported by LFA in 6.18.1:
  - MHOP IP BFD
  - RSVP—LFA should not be enabled if RSVP is being used.
  - BGP—LFA should not be enabled if BGP is being used.

## — TWAMP

- The SAOS assumes LDP is enabled before enabling LFA.
- The IP and LDP LFA functionality is inter-related, so it is mandatory to have LDP enabled on the IP interface on which IGP is enabled, for LFA functionality to work properly. The LFA functionality is only supported with LDP working with a PHP-enabled in network. If LDP is configured with the label type as non-reserved on some node, then LFA functionality may not work properly and such configuration (without PHP) is not supported.
- In 6.x devices when LFA is enabled, it internally disables the ECMP load balancing.
- The LFA functionality supports MPLS OAM for LDP LSPs.

In the deployment, IP bidirectional forwarding detection (BFD) is used to monitor the IP links and IP adjacency. When a fault occurs, the IP BFD triggers a next hop fault notification, which is used as a trigger to activate the backup path in the IP LFA functionality on the local source node. The IGP converges according to its timers, then a new primary path is calculated and programmed in the hardware, and then traffic is shifted to the new path. There may be minimal traffic affected during the new primary path programming.

## Hardware support

LFA is supported on all SAOS 6.x devices that support LDP. LFA performance depends on the hardware-assisted switchover capability of the device. The scalability of LFA is according to available hardware resources.

On 3926m and 3928 platforms, there may be slightly higher traffic drops compared with other platforms during new primary path programming when IGP converges.

This table shows the various SAOS 6.x devices that support LFA hardware-assisted switchovers.

**Table 22-1**  
**SAOS 6.x devices that support LFA hardware-assisted switchovers**

| Device                 | MPLS LDP LFA | IP LFA |
|------------------------|--------------|--------|
| 3916, 3930, 3931, 3932 | No           | No     |
| 3942, 5142, 5150, 5160 | Yes          | No     |
| 3926m, 3928            | Yes          | No     |

## Notes

- IP and LDP FRR support have base key functionality with limited scale.

- The IP FRR is software-based switching. There is no HW assist for IP traffic.
- If LFA is enabled on an IP interface, the IP interface can act as both a protected and a protector interface. By default, all point-to-point interfaces become LFA-enabled when global LFA is enabled.
- Node protection is enabled by default.
- Tie-breaker configuration is not provided. By default, node protection is prioritized over link protection.
- While LFA can be enabled/disabled per link, it is enabled or disabled for both protected and protector roles.
- Route-map based FRR functionality control is supported.
- To support IP/LDP FRR for IS-IS IGP, LFA works independently in each contiguous IS-IS level.
- LFA protection can be provided for inter-area prefixes in an IS-IS level, provided the primary and alternate next hop are within the same IS-IS level as the power-aware localized routing (PLR) protocol.
- Hardware-assisted FRR is possible only on devices that can support hardware-based MPLS switchover. For other devices, FRR is fully software-driven.
- LFA Enable or LFA Disable on a node should be done under maintenance window because they may impact the IP and LDP traffic transiting the node. It is recommended that you should take care that LFA Disable or Enable operations are done on individual nodes one by one and not simultaneously on multiple nodes of the network. You should also an interval of a few minutes between LFA Disable / Enable CLI execution on a node.
- You set an IP interface as point-to-point. (For instructions, see [Setting or unsetting the IP interface as a point-to-point interface](#).) Then all protocols running on that IP interface will be point-to-point. Alternatively, you can specify that the ISIS protocol on an IP interface should be point-to-point. (For instructions, see [Setting or unsetting the network type of the ISIS protocol on an IP interface](#)). Note that LFA is only supported on Point to Point interface with IS-IS as IGP.
- The LFA configuration of a max-delay timer value to calculate an alternate protection path is global and not per protocol.
- LFA disable disables LFA functionality and flushes existing LFA backups.
- When LFA is enabled, ECMP load balancing is internally disabled on the device.
- The micro-loop prevention timer configuration (as described in [RFC 8333](#)) is supported for IS-IS.

## Procedures

Procedures to configure and manage LFA include:

- [“Enabling or disabling global LFA” on page 22-14](#)
- [“Setting or unsetting LFA on a protocol or level” on page 22-15](#)
- [“Enabling or disabling LFA on an IP interface” on page 22-16](#)
- [“Setting or unsetting LFA on a route map” on page 22-17](#)
- [“Setting or unsetting the LFA calculation delay timer after IGP convergence” on page 22-18](#)
- [“Setting or unsetting the reserved IP address for LFA” on page 22-19](#)
- [“Setting the local micro-loop prevention delay timer for IS-IS” on page 22-20](#)
- [“Setting or unsetting the IP interface as a point-to-point interface” on page 22-22](#)
- [“Managing route maps” on page 22-23](#)
- [“Managing prefix lists” on page 22-24](#)
- [“Displaying LFA information” on page 22-25](#)

## Procedure 22-1

### Enabling or disabling global LFA

---

Enable or disable global LFA functionality.

---

| Step | Action |
|------|--------|
|------|--------|

---

- |   |                                                                                                                                                                                                            |
|---|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | <p>Enable global LFA functionality:</p> <pre>ip fast-reroute &lt;enable   disable&gt;</pre> <p>where</p> <p>&lt;enable  <br/>disable&gt; is the action taken for global LFA functionality</p> <p>—end—</p> |
|---|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## Procedure 22-2

# Setting or unsetting LFA on a protocol or level

Set or unset LFA on a protocol or IS-IS level.

---

## StepAction

---

### *To set LFA on a protocol*

- 1 Set LFA on a protocol:  

```
ip fast-reroute <set | unset> [src-proto <isis>]
```

where

|               |                                     |
|---------------|-------------------------------------|
| <set   unset> | is the action taken on the protocol |
| isis          | is the name of the protocol.        |

### *To set LFA on an IS-IS level*

- 2 Set LFA on an IS-IS level:  

```
ip fast-reroute <set | unset> [src-proto <isis> level <L1
| L2>]
```

where

|               |                                        |
|---------------|----------------------------------------|
| <set   unset> | is the action taken on the IS-IS level |
| <L1   L2>     | is the IS-IS level                     |

—end—



## Procedure 22-3

### Enabling or disabling LFA on an IP interface

---

Enable or disable LFA functionality on an IP interface.

If LFA is disabled on an IP interface, the IP interface will not act as protector for any prefix, and the IP interface will not be protected by any other NHOP.

---

| Step | Action |
|------|--------|
|------|--------|

---

|   |                                |
|---|--------------------------------|
| 1 | Enable LFA on an IP interface: |
|---|--------------------------------|

```
ip fast-reroute interface <enable | disable> ip-interface
<ip-interface>
```

where

<enable |  
disable> is the action taken for LFA

<ip-interface> is the name of the IP interface

—end—

---

## Procedure 22-4

### Setting or unsetting LFA on a route map

---

Set or unset LFA functionality on a route map.

| Step             | Action                                                                                                                                                                                                                                                                                                                     |               |                                       |                  |                              |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|---------------------------------------|------------------|------------------------------|
| 1                | <p>Set LFA on a route map:</p> <pre>ip fast-reroute &lt;set   unset&gt; [route-map &lt;route-map-name&gt;]</pre> <p>where</p> <table><tr><td>&lt;set   unset&gt;</td><td>is the action taken for the route map</td></tr><tr><td>&lt;route-map-name&gt;</td><td>is the name of the route map</td></tr></table> <p>—end—</p> | <set   unset> | is the action taken for the route map | <route-map-name> | is the name of the route map |
| <set   unset>    | is the action taken for the route map                                                                                                                                                                                                                                                                                      |               |                                       |                  |                              |
| <route-map-name> | is the name of the route map                                                                                                                                                                                                                                                                                               |               |                                       |                  |                              |

## Procedure 22-5

### Setting or unsetting the LFA calculation delay timer after IGP convergence

---

Set or unset the LFA calculation delay timer after IGP convergence due to IGP or topology changes.

**Note:** The delay timer may reset after any topology change.

---

| Step | Action |
|------|--------|
|------|--------|

---

- |   |                                                            |
|---|------------------------------------------------------------|
| 1 | Set the LFA calculation delay timer after IGP convergence: |
|---|------------------------------------------------------------|

```
ip fast-reroute <set | unset> [lfa-calc-max-delay <timer-
value-in-msec>]
```

where

<set | unset> is the action taken for the delay timer

<timer-value-in-msec> is the number of milliseconds to delay the LFA calculation after IGP convergence

—end—

---

## Procedure 22-6

### Setting or unsetting the reserved IP address for LFA

---

Set or unset the reserved IP address for LFA internal functionality.

---

| Step | Action |
|------|--------|
|------|--------|

---

- |   |                                                                                                                          |
|---|--------------------------------------------------------------------------------------------------------------------------|
| 1 | <p>Set the reserved IP address for LFA:</p> <pre>ip fast-reroute set lfa-reserved-ip 192.168.100.2/31</pre> <p>—end—</p> |
|---|--------------------------------------------------------------------------------------------------------------------------|

## Procedure 22-7

# Setting the local micro-loop prevention delay timer for IS-IS

---

Set the local micro-loop prevention delay timer for IS-IS.

---

| Step | Action |
|------|--------|
|------|--------|

---

- |                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                 |                                   |                         |                                                                                                                                                        |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|-----------------------------------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1                       | <p>Set the local micro-loop prevention delay timer for IS-IS:</p> <pre>isis instance set isis-instance &lt;isis-instance&gt; uloop-<br/>prevention-timer &lt;timer-value (msec)&gt;</pre> <p>where</p> <table><tbody><tr><td>&lt;isis-instance&gt;</td><td>is the name of the IS-IS instance</td></tr><tr><td>&lt;timer-value<br/>(msec)&gt;</td><td>is the value by which to delay the route configuration<br/>on the local node after any link-down detection, so<br/>that local micro-loops will not occur</td></tr></tbody></table> <p>—end—</p> | <isis-instance> | is the name of the IS-IS instance | <timer-value<br>(msec)> | is the value by which to delay the route configuration<br>on the local node after any link-down detection, so<br>that local micro-loops will not occur |
| <isis-instance>         | is the name of the IS-IS instance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                 |                                   |                         |                                                                                                                                                        |
| <timer-value<br>(msec)> | is the value by which to delay the route configuration<br>on the local node after any link-down detection, so<br>that local micro-loops will not occur                                                                                                                                                                                                                                                                                                                                                                                               |                 |                                   |                         |                                                                                                                                                        |

## Procedure 22-8

# Setting or unsetting the network type of the ISIS protocol on an IP interface

Use this procedure to set or unset the network type (point-to-point or broadcast) of the ISIS protocol on an IP interface.

| Step | Action |
|------|--------|
|------|--------|

### *To set the network type of the ISIS protocol on an IP interface*

- 1 Set the network type of the ISIS protocol on an IP interface:  

```
isis interface set ip-interface <ip_interface>
network-type <point-to-point/broadcast>
```

where

<ip\_interface> is the IP interface name

<point-to-point  
|Broadcast> is the type of network

### *To unset the network type of the ISIS protocol on an IP interface*

- 2 Unset the network type of the ISIS protocol on an IP interface:  

```
isis interface unset set ip-interface <ip_interface>
network-type <point-to-point/broadcast>
```

where

<ip\_interface> is the IP interface name

<point-to-point  
|Broadcast> is the type of network

—end—

## Example

This example shows how to set the network type of the ISIS protocol on an IP interface to point-to-point.

```
isis interface set ip-interface myInterface network-type
point-to-point
```

## Procedure 22-9

### Setting or unsetting the IP interface as a point-to-point interface

---

Set or unset the IP interface as a point-to-point interface.

---

| Step | Action |
|------|--------|
|------|--------|

---

- |                                |                                                                                                                                                                                                                                                                                                                                                                                           |                |                          |                                |                        |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|--------------------------|--------------------------------|------------------------|
| 1                              | <p>Set the IP interface as a point-to-point interface:</p> <pre>interface set ip-interface &lt;ip_interface&gt; network-type &lt;Broadcast point-to-point&gt;</pre> <p>where</p> <table><tbody><tr><td>&lt;ip_interface&gt;</td><td>is the IP interface name</td></tr><tr><td>&lt;Broadcast<br/> point-to-point&gt;</td><td>is the type of network</td></tr></tbody></table> <p>—end—</p> | <ip_interface> | is the IP interface name | <Broadcast<br> point-to-point> | is the type of network |
| <ip_interface>                 | is the IP interface name                                                                                                                                                                                                                                                                                                                                                                  |                |                          |                                |                        |
| <Broadcast<br> point-to-point> | is the type of network                                                                                                                                                                                                                                                                                                                                                                    |                |                          |                                |                        |

### Example

This example shows how to set the IP interface as a point-to-point interface.

```
interface set ip-interface myInterface network-type
point-to-point
```

## Procedure 22-10

### Managing route maps

Create and manage route maps.

| Step | Action |
|------|--------|
|------|--------|

#### *To create a route map*

- 1 Create a route map:  

```
route-map create map-name <map-name>
```

where  
 <map-name> is the name of the route map

#### *To delete a route map*

- 2 Delete a route map:  

```
route-map delete map-name <map-name>
```

#### *To remove an instance of a route map*

- 3 Remove an instance of a route map:  

```
route-map remove map-name <map-name> [seq <NUMBER: 10-65535>]
```

where  
 [seq <NUMBER: 10-65535>] is the sequence number of the route map instance

#### *To add an instance of a route map for a specific IP address, action, and prefix list*

- 4 Add an instance of a route map for a specific IP address, action, and prefix list:  

```
route-map match-address add map-name <map-name> [seq <NUMBER: 10-65535>] [action <permit | deny>] [prefix-list-name <list-name>]
```

where  
 [action <permit | deny>] is the action taken by the device if the route map instance matches the IP address  
 [prefix-list-name <list-name>] is the name of the prefix list

#### *To remove an instance of a route map for a specific IP address and action*

- 5 Remove an instance of a route map for a specific IP address and action:  

```
route-map match-address remove map-name <map-name> [seq <NUMBER: 10-65535>] [action <permit | deny>]
```

—end—



## Procedure 22-11

### Managing prefix lists

Create and manage prefix lists.

| Step | Action |
|------|--------|
|------|--------|

#### *To create a prefix list*

- 1 Create a prefix list:  

```
prefix-list create list-name <list-name>
```

where  
<list-name> is the name of the prefix list

#### *To delete a prefix list*

- 2 Delete a prefix list:  

```
prefix-list delete list-name <list-name>
```

#### *To add an instance of a prefix list*

- 3 Add an instance of a prefix list:  

```
prefix-list add list-name <list-name> [seq <NUMBER: 1-4294967295>] [action <permit | deny>] [ipaddress/mask <IP address in CIDR notation>] [ge <NUMBER: 1-32>] [le <NUMBER: 1-32>]
```

where

|                                              |                                                                                                                            |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| [seq <NUMBER: 1-4294967295>]                 | is the sequence number of the prefix list instance                                                                         |
| [action <permit   deny>]                     | is the action taken by the device if the prefix list instance matches an IP address/mask, minimum or maximum prefix length |
| ipaddress/mask <IP address in CIDR notation> | is the ip-address with subnet mask                                                                                         |
| [ge <NUMBER: 1-32>]                          | is the minimum prefix length to be matched                                                                                 |
| [le <NUMBER: 1-32>]                          | is the maximum prefix length to be matched                                                                                 |

#### *To remove an instance of a prefix list*

- 4 Remove an instance of a prefix list:  

```
prefix-list remove list-name <list-name> [seq <NUMBER: 1-4294967295>]
```

—end—

## Procedure 22-12

### Displaying LFA information

Display LFA information, including summary statistics, interface, path, and route details.

| Step | Action |
|------|--------|
|------|--------|

#### *To display LFA information*

- 1 Display LFA information:  
`ip fast-reroute show`

#### *To display a summary of LFA statistics*

- 2 Display a summary of LFA statistics:  
`ip fast-reroute show lfa-summary`

#### *To display LFA interfaces*

- 3 Display LFA interfaces:  
`ip fast-reroute interface show`

#### *To display LFA paths*

- 4 Display LFA paths:  
`ip fast-reroute show lfa-path-table`

#### *To display LFA paths per prefix*

- 5 Display LFA paths per prefix:  
`ip fast-reroute show lfa-path-table [prefix <prefix-ip>]`  
where  
`<prefix-ip>` is the IP address of the prefix

#### *To display routing entries, including LFA routes*

- 6 Display routing entries:  
`ip route show`

#### *To display LFA route changes*

- 7 Display LFA route changes:  
`ip fib show`

#### *To display the interfaces attached to an ISIS instance*

- 8 Display the interfaces attached to an ISIS instance:  
`isis instance show isis-instance <isis-instance>`  
`attached-interfaces`

—end—



---

## Sample MPLS topology

---

This chapter describes the following sample MPLS topologies and provides the commands to configure the sample topologies.


- “H-VPLS configuration example mixed platform” on page 23-2
- “VPLS with CFM configuration example 3916 and 3930” on page 23-9
- “G.8032 and VPLS interoperability example” on page 23-13
- “MPLS-TP configuration example” on page 23-18
- “End-to-End EPL service over MPLS example” on page 23-20
- “Resource affinity configuration for a TE interface example” on page 23-23
- “Encapsulation configuration examples” on page 23-25

## Procedure 23-1

### H-VPLS configuration example mixed platform

This example provides a mix of 39XX/51XX and 5410 devices with a VPLS core and spoke virtual circuits as shown in [Figure 23-1](#).

**Figure 23-1**  
H-VPLS topology mixed platform



## H-VPLS configuration example on 5410

The following sample configuration file segment shows the H-VPLS configuration for the 5410 in [Figure 23-1 on page 23-2](#).

```

rstp disable
mstp disable
aggregation disable
system shell set global-inactivity-timer off

port disable port 6/1-6/4
port enable port 6/4
port enable port 6/3

virtual-switch create vs test-vs

cpu-interface sub-interface create cpu-subinterface rg
cpu-interface sub-interface show

sub-port create sub-port sp6/4_1 parent-port 6/4 classifier-precedence 1
egress-l2-transform push-8100.305.1
sub-port add sub-port sp6/4_1 class-element 1 vtag-stack 305
sub-port create sub-port sp6/3_1 parent-port 6/3 classifier-precedence 1
egress-l2-transform push-8100.10.1
sub-port add sub-port sp6/3_1 class-element 1 vtag-stack 10

virtual-switch interface attach sub-port sp6/4_1 vs test-vs
virtual-switch interface attach cpu-subinterface rg vs test-vs

interface create loopback test-loop ip 3.100.1.1
interface create ip-interface test-ip ip 4.100.160.
100 subnet 255.255.255.0 vs test-vs

rsvp-te enable ip-interface test-loop
rsvp-te enable ip-interface test-ip
rsvp-te enable

ospf instance create ospf-instance ospfInst1
ospf instance enable ospf-instance ospfInst1
ospf instance set ospf-instance ospfInst1 router-id 3.100.1.1
ospf interface attach ip-interface test-ip ospf-instance ospfInst1
ospf interface enable ip-interface test-ip
ospf interface attach ip-interface test-loop ospf-instance ospfInst1
ospf interface enable ip-interface test-loop

mpls tunnel create rsvp-ingress dyn-100_160 dest-ip 3.160.1.1
mpls l2-vpn create dynamic-vc dynvc-100_160 peer 3.160.1.1 tunnel dyn-100_160
pw-id 109

ldp enable

virtual-switch create vs vpls
virtual-switch interface attach sub-port sp6/3_1 vs vpls
virtual-switch interface attach mpls-vc dynvc-100_160 vs vpls

```

## H-VPLS configuration example on 3916

The following sample configuration file segment shows the H-VPLS configuration for the 3916 in [Figure 23-1 on page 23-2](#).

```

rstp disable
mstp disable
aggregation disable
system shell set global-inactivity-timer off

port disable port 1-6
port enable port 4
port enable port 2
port enable port 5
port enable port 6

vlan create vlan 301
vlan create vlan 302
vlan create vlan 304

vlan add vlan 301 port 2
vlan add vlan 302 port 4
vlan add vlan 304 port 5

interface create loopback loopback-B ip 3.116.1.1
interface create ip-interface intf-116_150 ip 4.116.150.116/24 vlan 302 ip-
forwarding on
interface create ip-interface intf-116_131 ip 4.116.131.116/24 vlan 301 ip-
forwarding on
interface create ip-interface intf-116_160 ip 4.116.160.116/24 vlan 304 ip-
forwarding on

ospf instance create ospf-instance OSPF-Inst-B
ospf instance set ospf-instance OSPF-Inst-B router-id 3.116.1.1
ospf instance enable ospf-instance OSPF-Inst-B

ospf interface attach ip-interface loopback-B ospf-instance OSPF-Inst-B
ospf interface enable ip-interface loopback-B
ospf interface attach ip-interface intf-116_150 ospf-instance OSPF-Inst-B
ospf interface enable ip-interface intf-116_150
ospf interface attach ip-interface intf-116_131 ospf-instance OSPF-Inst-B
ospf interface enable ip-interface intf-116_131
ospf interface attach ip-interface intf-116_160 ospf-instance OSPF-Inst-B
ospf interface enable ip-interface intf-116_160

rsvp-te set ip-interface loopback-B advertised-label non-reserved
rsvp-te set ip-interface intf-116_150 advertised-label implicit-null
rsvp-te set ip-interface intf-116_131 advertised-label non-reserved
rsvp-te set ip-interface intf-116_160 advertised-label non-reserved
rsvp-te enable ip-interface loopback-B
rsvp-te enable ip-interface intf-116_150
rsvp-te enable ip-interface intf-116_131
rsvp-te enable ip-interface intf-116_160
rsvp-te enable

mpls tunnel create rsvp-ingress rsvp-itnl-116_131 dest-ip 3.131.1.1
mpls tunnel create rsvp-ingress rsvp-itnl-116_160 dest-ip 3.160.1.1

```

```

mpls tunnel create static-ingress stat-itnl-116_150 dest-ip 3.150.1.1 next-
hop-ip 4.116.150.150 label
105
mpls tunnel create static-egress stat-etnl-116_150
label 106

ldp enable

mpls l2-vpn create dynamic-vc dyn-116_131 pw-id 105 peer 3.131.1.1 tunnel
rsvp-itnl-116_131 pw-mode mesh
mpls l2-vpn create dynamic-vc dyn-116_160 pw-id 102 peer 3.160.1.1 tunnel
rsvp-itnl-116_160 pw-mode mesh
mpls l2-vpn create dynamic-vc dyn-116_150 pw-id 103 peer 3.150.1.1 tunnel
stat-itnl-116_150 pw-mode spoke

virtual-switch ethernet create vs vs-A mode vpls
virtual-switch ethernet attach mpls-vc dyn-116_131 vs vs-A
virtual-switch ethernet attach mpls-vc dyn-116_160 vs vs-A
virtual-switch ethernet attach mpls-vc dyn-116_150 vs vs-A

```

## H-VPLS configuration example on 3930

The following sample configuration file segment shows the H-VPLS configuration for the 3930 in [Figure 23-1 on page 23-2](#).

```

rstp disable
mstp disable
aggregation disable
system shell set global-inactivity-timer off

port disable port 1-10

port enable port 1
port enable port 9

vlan create vlan 300
vlan create vlan 10

vlan add vlan 300 port 9
vlan add vlan 10 port 1

interface create loopback loopback-A ip 3.130.1.1
interface create ip-interface intf-130_131 ip 4.130.131.130/24 vlan 300 ip-
forwarding on

ospf instance create ospf-instance OSPF-Inst-A
ospf instance set ospf-instance OSPF-Inst-A router-id 3.130.1.1
ospf instance enable ospf-instance OSPF-Inst-A
ospf interface attach ip-interface loopback-A ospf-instance OSPF-Inst-A
ospf interface enable ip-interface loopback-A
ospf interface attach ip-interface intf-130_131 ospf-instance OSPF-Inst-A
ospf interface enable ip-interface intf-130_131

rsvp-te set ip-interface loopback-A advertised-label non-reserved
rsvp-te set ip-interface intf-130_131 advertised-label implicit-null
rsvp-te enable ip-interface loopback-A
rsvp-te enable ip-interface intf-130_131
rsvp-te enable

```



## 23-6 Sample MPLS topology

---

```
mpls tunnel create rsvp-ingress rsvp-itnl-130_131 dest-ip 3.131.1.1
ldp enable

mpls l2-vpn create dynamic-vc dyn-130_131 pw-id 100 peer 3.131.1.1 tunnel
rsvp-itnl-130_131 pw-mode spoke

virtual-switch ethernet create vs vs-A mode vpws
virtual-switch ethernet add vs vs-A port 1 vlan 10
virtual-switch ethernet attach mpls-vc dyn-130_131 vs vs-A
```

### H-VPLS configuration example on 3931

The following sample configuration file segment shows the H-VPLS configuration for the 3931 in [Figure 23-1 on page 23-2](#).

```
rstp disable
mstp disable
aggregation disable
system shell set global-inactivity-timer off

port disable port 1-10
port enable port 9
port enable port 5
port enable port 1
port enable port 6
port set port 5 mirror-port on

vlan create vlan 300
vlan create vlan 301
vlan create vlan 303

vlan add vlan 300 port 9
vlan add vlan 301 port 6
vlan add vlan 303 port 1

interface create loopback loopback-C ip 3.131.1.1
interface create ip-interface intf-131_130 ip 4.130131.131/24 vlan 300 ip-
forwarding on
interface create ip-interface intf-131_116 ip 4.116131.131/24 vlan 301 ip-
forwarding on
interface create ip-interface intf-131_160 ip 4.131160.131/24 vlan 303 ip-
forwarding on

ospf instance create ospf-instance OSPF-Inst-C
ospf instance set ospf-instance OSPF-Inst-C router-id 3.131.1.1
ospf instance enable ospf-instance OSPF-Inst-C
ospf interface attach ip-interface loopback-C ospf-instance OSPF-Inst-C
ospf interface enable ip-interface loopback-C
ospf interface attach ip-interface intf-131_130 ospf-instance OSPF-Inst-C
ospf interface enable ip-interface intf-131_130
ospf interface attach ip-interface intf-131_116 ospf-instance OSPF-Inst-C
ospf interface enable ip-interface intf-131_116
ospf interface attach ip-interface intf-131_160 ospf-instance OSPF-Inst-C
ospf interface enable ip-interface intf-131_160

rsvp-te set ip-interface loopback-C advertised-label non-reserved
rsvp-te set ip-interface intf-131_130 advertised-label non-reserved
rsvp-te set ip-interface intf-131_116 advertised-label non-reserved
```

```

rsvp-te set ip-interface intf-131_160 advertised-label non-reserved
rsvp-te enable ip-interface loopback-C
rsvp-te enable ip-interface intf-131_130
rsvp-te enable ip-interface intf-131_116
rsvp-te enable ip-interface intf-131_160
rsvp-te enable

mpls tunnel create rsvp-ingress rsvp-itnl-131_130 dest-ip 3.130.1.1
mpls tunnel create rsvp-ingress rsvp-itnl-131_116 dest-ip 3.116.1.1
mpls tunnel create rsvp-ingress rsvp-itnl-131_160 dest-ip 3.160.1.1

ldp enable

mpls l2-vpn create dynamic-vc dyn-131_130 pw-id 100 peer 3.130.1.1 tunnel
rsvp-itnl-131_130 pw-mode spoke
mpls l2-vpn create dynamic-vc dyn-131_116 pw-id 105 peer 3.116.1.1 tunnel
rsvp-itnl-131_116 pw-mode mesh
mpls l2-vpn create dynamic-vc dyn-131_160 pw-id 106 peer 3.160.1.1 tunnel
rsvp-itnl-131_160 pw-mode mesh

virtual-switch ethernet create vs vs-A mode vpls
virtual-switch ethernet attach mpls-vc dyn-131_130 vs vs-A
virtual-switch ethernet attach mpls-vc dyn-131_116 vs vs-A
virtual-switch ethernet attach mpls-vc dyn-131_160 vs vs-A
virtual-switch ethernet add vs vs-A port 5 vlan 15
virtual-switch ethernet attach mpls-vc dyn-130_131 vs vs-A

```

## H-VPLS configuration example on 5150

The following sample configuration file segment shows the H-VPLS configuration for the 5150 in [Figure 23-1 on page 23-2](#).

```

rstp disable
mstp disable
aggregation disable

system shell set global-inactivity-timer off

port disable port 1.1-1.48
port disable port 2.1-2.2
port disable port 3.1-3.2
port enable port 1.8
port enable port 1.1

vlan create vlan 302
vlan create vlan 10
vlan add vlan 302 port 1.1
vlan add vlan 10 port 1.8

interface create loopback loopback-D ip 3.150.1.1
interface create ip-interface intf-150_116 ip 4.116.150.150/24 vlan 302 ip-
forwarding on

ospf instance create ospf-instance OSPF-Inst-D
ospf instance set ospf-instance OSPF-Inst-D router-id 3.150.1.1
ospf instance enable ospf-instance OSPF-Inst-D
ospf interface attach ip-interface loopback-D ospf-instance OSPF-Inst-D
ospf interface enable ip-interface loopback-D
ospf interface attach ip-interface intf-150_116 ospf-instance OSPF-Inst-D

```

## 23-8 Sample MPLS topology

---

```
ospf interface enable ip-interface intf-150_116

rsvp-te set ip-interface loopback-D advertised-label non-reserved
rsvp-te set ip-interface intf-150_116 advertised-label non-reserved
rsvp-te enable ip-interface loopback-D
rsvp-te enable ip-interface intf-150_116
rsvp-te enable

mpls tunnel create static-ingress stat-itnl-150_116 dest-ip 3.116.1.1 next-
hop-ip 4.116.150.116 label 106
mpls tunnel create static-egress stat-etnl-150_116 label 105

ldp enable

mpls l2-vpn create dynamic-vc dyn-150_116 pw-id 103 peer 3.116.1.1 tunnel
stat-itnl-150_116 pw-mode spoke


virtual-switch ethernet create vs vs-A mode vpws
virtual-switch ethernet add vs vs-A port 1.8 vlan 10
virtual-switch ethernet attach mpls-vc dyn-150_116 vs vs-A
```

## Procedure 23-2

### VPLS with CFM configuration example 3916 and 3930

This example shows a VPLS between two 3916 switches and a 3930 as shown in [Figure 23-2](#) with dynamic virtual circuits. The CFM service is configured over the data virtual switch and virtual circuit endpoints have down MEPs.

**Figure 23-2**  
VPLS with CFM topology



### VPLS configuration example on 3916-80

This example shows the VPLS configuration of 3916-80 shown in [Figure 23-2](#).

```

system set host-name CN3916-80
rstp disable
mstp disable
aggregation disable
system shell set global-inactivity-timer off

port disable port 1,5,6
port enable port 2,3,4

vlan create vlan 200
vlan create vlan 301
vlan create vlan 302

vlan add vlan 200 port 3
vlan add vlan 301 port 2
vlan add vlan 302 port 4

interface create loopback LBK ip 2.2.2.2
interface create ip-interface IFACE-4.1.1.2 ip 4.1.1.2/24 vlan 301 ip-
forwarding on

```

## 23-10 Sample MPLS topology

---

```
interface create ip-interface IFACE-6.1.1.1 ip 6.1.1.1/24 vlan 302 ip-
forwarding on

ospf instance create ospf-instance OSPF-INST
ospf instance set ospf-instance OSPF-INST router-id 2.2.2.2
ospf instance enable ospf-instance OSPF-INST

ospf interface attach ip-interface LBK ospf-instance OSPF-INST
ospf interface enable ip-interface LBK
ospf interface attach ip-interface IFACE-4.1.1.2 ospf-instance OSPF-INST
ospf interface enable ip-interface IFACE-4.1.1.2
ospf interface attach ip-interface IFACE-6.1.1.1 ospf-instance OSPF-INST
ospf interface enable ip-interface IFACE-6.1.1.1

rsvp-te enable
rsvp-te set ip-interface LBK advertised-label non-reserved
rsvp-te set ip-interface IFACE-4.1.1.2 advertised-label non-reserved
rsvp-te set ip-interface IFACE-6.1.1.1 advertised-label non-reserved
rsvp-te enable ip-interface LBK
rsvp-te enable ip-interface IFACE-4.1.1.2
rsvp-te enable ip-interface IFACE-6.1.1.1

mpls dynamic-label-range set min-label 200 max-label 299
mpls tunnel create rsvp-ingress rsvp-itnl-4.1.1 dest-ip 1.1.1.1
mpls tunnel create rsvp-ingress rsvp-itnl-6.1.1 dest-ip 3.3.3.3

ldp enable

mpls l2-vpn create dynamic-vc dyn-4.1.1 pw-id 105 peer 1.1.1.1 tunnel rsvp-
itnl-4.1.1 pw-mode mesh
mpls l2-vpn create dynamic-vc dyn-6.1.1 pw-id 102 peer 3.3.3.3 tunnel rsvp-
itnl-6.1.1 pw-mode mesh

virtual-switch ethernet create vs vs-A mode vpls
virtual-switch ethernet add vs vs-A port 3 vlan 200
virtual-switch ethernet attach mpls-vc dyn-4.1.1 vs vs-A
virtual-switch ethernet attach mpls-vc dyn-6.1.1 vs vs-A
```

## VPLS configuration example on 3916-60

This example shows the VPLS configuration of 3916-60 shown in [Figure 23-2 on page 23-9](#).

```
system set host-name CN3916-60
rstp disable
mstp disable
aggregation disable
system shell set global-inactivity-timer off

port disable port 3,4,5
port enable port 1,2,6

vlan create vlan 100
vlan create vlan 301
vlan create vlan 303

vlan add vlan 301 port 2
vlan add vlan 303 port 6
```

```

interface create loopback LBK ip 1.1.1.1
interface create ip-interface IFACE-4.1.1.1 ip 4.1.1.1/24 vlan 301 ip-
forwarding on
interface create ip-interface IFACE-5.1.1.2 ip 5.1.1.2/24 vlan 303 ip-
forwarding on

ospf instance create ospf-instance OSPF-INST
ospf instance set ospf-instance OSPF-INST router-id 1.1.1.1
ospf instance enable ospf-instance OSPF-INST

ospf interface attach ip-interface LBK ospf-instance OSPF-INST
ospf interface enable ip-interface LBK
ospf interface attach ip-interface IFACE-4.1.1.1 ospf-instance OSPF-INST
ospf interface enable ip-interface IFACE-4.1.1.1
ospf interface attach ip-interface IFACE-5.1.1.2 ospf-instance OSPF-INST
ospf interface enable ip-interface IFACE-5.1.1.2

rsvp-te set ip-interface LBK advertised-label non-reserved
rsvp-te set ip-interface IFACE-4.1.1.1 advertised-label non-reserved
rsvp-te set ip-interface IFACE-5.1.1.2 advertised-label non-reserved

rsvp-te enable
rsvp-te enable ip-interface LBK
rsvp-te enable ip-interface IFACE-4.1.1.1
rsvp-te enable ip-interface IFACE-5.1.1.2

mpls dynamic-label-range set min-label 400 max-label 499
mpls tunnel create rsvp-ingress rsvp-itnl-4.1.1 dest-ip 2.2.2.2
mpls tunnel create rsvp-ingress rsvp-itnl-5.1.1 dest-ip 3.3.3.3

ldp enable

mpls l2-vpn create dynamic-vc dyn-4.1.1 pw-id 105 peer 2.2.2.2 tunnel rsvp-
itnl-4.1.1 pw-mode mesh
mpls l2-vpn create dynamic-vc dyn-5.1.1 pw-id 106 peer 3.3.3.3 tunnel rsvp-
itnl-5.1.1 pw-mode mesh

virtual-switch ethernet create vs vs-A mode vpls
virtual-switch ethernet attach mpls-vc dyn-4.1.1 vs vs-A
virtual-switch ethernet attach mpls-vc dyn-5.1.1 vs vs-A

virtual-switch ethernet add vs vs-A port 1 vlan 100

```

## VPLS configuration example on 3930-14

This example shows the VPLS configuration of 3930-14 shown in [Figure 23-2 on page 23-9](#).

```

system set host-name CN3930-14
rstp disable
mstp disable
aggregation disable
system shell set global-inactivity-timer off

port disable port 1,2,3,5,7,8,9,10,11,12
port enable port 4,6

vlan create vlan 302
vlan create vlan 303

```

## 23-12 Sample MPLS topology

---

```
vlan add vlan 303 port 6
vlan add vlan 302 port 4

interface create loopback LBK ip 3.3.3.3
interface create ip-interface IFACE-5.1.1.1 ip 5.1.1.1/24 vlan 303 ip-
forwarding on
interface create ip-interface IFACE-6.1.1.2 ip 6.1.1.2/24 vlan 302 ip-
forwarding on

ospf instance create ospf-instance OSPF-INST
ospf instance set ospf-instance OSPF-INST router-id 3.3.3.3
ospf instance enable ospf-instance OSPF-INST
ospf interface attach ip-interface LBK ospf-instance OSPF-INST
ospf interface enable ip-interface LBK
ospf interface attach ip-interface IFACE-5.1.1.1 ospf-instance OSPF-INST
ospf interface enable ip-interface IFACE-5.1.1.1
ospf interface attach ip-interface IFACE-6.1.1.2 ospf-instance OSPF-INST
ospf interface enable ip-interface IFACE-6.1.1.2

rsvp-te set ip-interface LBK advertised-label non-reserved
rsvp-te set ip-interface IFACE-5.1.1.1 advertised-label non-reserved
rsvp-te set ip-interface IFACE-6.1.1.2 advertised-label non-reserved
rsvp-te enable ip-interface LBK
rsvp-te enable ip-interface IFACE-5.1.1.1
rsvp-te enable ip-interface IFACE-6.1.1.2
rsvp-te enable

mpls dynamic-label-range set min-label 100 max-label 199
mpls tunnel create rsvp-ingress rsvp-itnl-5.1.1 dest-ip 1.1.1.1
mpls tunnel create rsvp-ingress rsvp-itnl-6.1.1 dest-ip 2.2.2.2

ldp enable

mpls l2-vpn create dynamic-vc dyn-5.1.1 pw-id 106 peer 1.1.1.1 tunnel rsvp-
itnl-5.1.1 pw-mode mesh
mpls l2-vpn create dynamic-vc dyn-6.1.1 pw-id 102 peer 2.2.2.2 tunnel rsvp-
itnl-6.1.1 pw-mode mesh


virtual-switch ethernet create vs vs-A mode vpls
virtual-switch ethernet attach mpls-vc dyn-5.1.1 vs vs-A
virtual-switch ethernet attach mpls-vc dyn-6.1.1 vs vs-A
```

## Procedure 23-3

### G.8032 and VPLS interoperability example

This example shows VPLS interoperability with G.8032 as shown in [Figure 23-3](#).

**Figure 23-3**  
G.8032 and VPLS topology



For the scenario where all of the following are true:

- virtual circuit is tagged
- attachment circuit on the right hand is port classification only (no service delimited tag)
- The tagged traffic is actually coming in over the attachment circuit, traffic that egresses the ring ports on the interconnect device are double-tagged.

The inner tag is the subscriber tag and the outer tag is the service delimiter VID pushed on by the virtual circuit. This is considered a mis-configuration and an acceptable result.

When only the second is true, VLAN 801 and VLAN 802 traffic will actually go out over the east and west (sub) ports that are in VS2, even though there is a VR on that device that is actually blocking one of those ports. That's considered an acceptable result and is also considered a misconfiguration.

### G.8032 and VPLS configuration example on 5410

This example shows the VPLS configuration of the 5410 shown in [Figure 23-3](#) on page 23-13.



## 23-14 Sample MPLS topology

---

```
virtual-switch create vs test-vs-prim
virtual-switch create vs vpls
virtual-switch create vs raps_vs
virtual-switch create vs raps-vs-900

system set host-name CN5410-3

interface create loopback test-loop ip 10.10.10.10
interface create ip-interface test-ip-prim ip 1.1.1.1 subnet 255.255.255.0 vs
test-vs-prim

cpu-interface sub-interface create cpu-subinterface test-cpu-prim

sub-port create sub-port sp6.3 parent-port 6/3 classifier-precedence 1 egress-
l2-transform push-8100.301.1
sub-port add sub-port sp6.3 class-element 1 vtag-stack 301
sub-port create sub-port sp6.6-800 parent-port 6/6 classifier-precedence 100
sub-port add sub-port sp6.6-800 class-element 1 vtag-stack 801
sub-port add sub-port sp6.6-800 class-element 2 vtag-stack 802
sub-port create sub-port sp6.4-800 parent-port 6/4 classifier-precedence 100
sub-port add sub-port sp6.4-800 class-element 1 vtag-stack 801
sub-port add sub-port sp6.4-800 class-element 2 vtag-stack 802
sub-port create sub-port sp6.6-900 parent-port 6/6 classifier-precedence 200
sub-port add sub-port sp6.6-900 class-element 1 vtag-stack 901
sub-port add sub-port sp6.6-900 class-element 2 vtag-stack 902
sub-port create sub-port sp6.4-900 parent-port 6/4 classifier-precedence 200
sub-port add sub-port sp6.4-900 class-element 1 vtag-stack 901
sub-port add sub-port sp6.4-900 class-element 2 vtag-stack 902

rsvp-te enable ip-interface test-loop
rsvp-te enable ip-interface test-ip-prim
rsvp-te path create rsvp-path path-prim logical-id 1
rsvp-te path set rsvp-path path-prim index 1 ip 1.1.1.2
rsvp-te enable

mpls tunnel create rsvp-ingress itnl-prim path path-prim dest-ip 30.30.30.30
logical-id 1

ldp enable

mpls l2-vpn create dynamic-vc test-dynVc pw-id 201
peer 30.30.30.30 tunnel itnl-prim pw-type eth-tagged

ospf instance create ospf-instance ospfInst1
ospf interface attach ip-interface test-loop ospf-instance ospfInst1
ospf interface enable ip-interface test-loop
ospf interface attach ip-interface test-ip-prim ospf-instance ospfInst1
ospf interface enable ip-interface test-ip-prim

virtual-switch interface attach sub-port sp6.3 vs test-vs-prim
virtual-switch interface attach cpu-subinterface test-cpu-prim vs test-vs-
prim
virtual-switch interface attach sub-port sp6.6-800 vs raps_vs
virtual-switch interface attach sub-port sp6.4-800 vs raps_vs
virtual-switch interface attach mpls-vc test-dynVc vs raps_vs

rstp disable

aggregation disable
```

```

ring-protection logical-ring create logical-ring-name LR1 ring-id 1 west-port
6/6 east-port 6/4 WTR 1
ring-protection virtual-ring create virtual-ring-name VR1 logical-ring LR1
raps-vid 800
ring-protection virtual-ring set ring VR1 east-port-rpl owner
ring-protection virtual-ring add ring VR1 vs raps_vs
ring-protection virtual-ring create virtual-ring-name VR2 logical-ring LR1
raps-vid 900
ring-protection virtual-ring add ring VR2 vs raps-vs-900

dhcp client enable

system shell set global-more off

```

## G.8032 and VPLS configuration example on 3916

This example shows the VPLS configuration of 3916-100 shown in [Figure 23-3 on page 23-13](#).

```

vlan create vlan 301

system set host-name CN3916-100

interface create loopback LBK ip 30.30.30.30
interface create ip-interface LSR-IFACE-1 ip 1.1.1.2/24 vlan 301 ip-forwarding
on

vlan add vlan 301 port 3

virtual-switch ethernet create vs vs1 mode vpls
virtual-switch ethernet add vs vs1 port 6 vlan 1

ospf instance create ospf-instance OSPF-Inst
ospf instance disable ospf-instance OSPF-Inst
ospf instance set ospf-instance OSPF-Inst router-id 30.30.30.30
ospf instance enable ospf-instance OSPF-Inst
ospf interface attach ip-interface LBK ospf-instance OSPF-Inst
ospf interface set ip-interface LBK priority 1
ospf interface enable ip-interface LBK
ospf interface attach ip-interface LSR-IFACE-1 ospf-instance OSPF-Inst
ospf interface set ip-interface LSR-IFACE-1 priority 1
ospf interface enable ip-interface LSR-IFACE-1

rsvp-te enable ip-interface LBK
rsvp-te enable ip-interface LSR-IFACE-1
rsvp-te path create rsvp-path path1 logical-id 1
rsvp-te path set rsvp-path path1 index 1 ip 1.1.1.1

rsvp-te enable

mpls tunnel create rsvp-ingress RSVP-ITNL-2 explicit-tunnel-path path1 dest-
ip 10.10.10.10 logical-id 1

ldp enable

mpls l2-vpn create dynamic-vc PWDYN-1 pw-id 201 peer 10.10.10.10 tunnel RSVP-
ITNL-2 pw-type eth-tagged

virtual-switch ethernet attach mpls-vc PWDYN-1 vs vs1

```

```
rstp disable

mstp disable

aggregation disable

system shell set global-more off
system shell set global-inactivity-timer off
```

### G.8032 and VPLS configuration example on 3916

This example shows the VPLS configuration of 3916-80 shown in [Figure 23-3 on page 23-13](#).

```
vlan create vlan 801-802,901-902

system set host-name CN3916-80

port disable port 2
port disable port 5
port disable port 6

vlan add vlan 801-802,901-902 port 1
vlan remove vlan 1 port 1
vlan add vlan 801-802 port 3
vlan remove vlan 1 port 3
vlan add vlan 801-802,901-902 port 4
vlan remove vlan 1 port 4

virtual-switch ethernet create vs vsA mode vpws
virtual-switch ethernet create vs vsB mode vpls
virtual-switch ethernet create vs vsC

ring-protection logical-ring create logical-ring-name LR1 ring-id 1 west-port
4 east-port 1 WTR 1
ring-protection virtual-ring create virtual-ring-name VR1 logical-ring LR1
raps-vid 800
ring-protection virtual-ring add ring VR1 vid 801-802
ring-protection virtual-ring create virtual-ring-name VR2 logical-ring LR1
raps-vid 900
ring-protection virtual-ring add ring VR2 vid 901-902

system shell set global-inactivity-timer off
```

### G.8032 and VPLS configuration example on 3916

This example shows the VPLS configuration of 3916-60 shown in [Figure 23-3 on page 23-13](#).

```
vlan create vlan 801-802,901-902

system set host-name CN3916-60

port disable port 3
port disable port 4
port disable port 5
```

```

vlan add vlan 801-802 port 1
vlan remove vlan 1 port 1
vlan add vlan 801-802,901-902 port 2
vlan remove vlan 1 port 2
vlan add vlan 801-802,901-902 port 6
vlan remove vlan 1 port 6

ring-protection logical-ring create logical-ring-name LR1 ring-id 1 west-port
2 east-port 6 WTR 1
ring-protection virtual-ring create virtual-ring-name VR1 logical-ring LR1
raps-vid 800
ring-protection virtual-ring add ring VR1 vid 801-802
ring-protection virtual-ring create virtual-ring-name VR2 logical-ring LR1
raps-vid 900
ring-protection virtual-ring add ring VR2 vid 901-902

system shell set global-inactivity-timer off

```

## G.8032 and VPLS configuration example on 3916

This example shows the VPLS configuration of 3916-14 shown in [Figure 23-3 on page 23-13](#).

```

vlan create vlan 801-802,901-902

system set host-name CN3916-14

vlan add vlan 801-802,901-902 port 1
vlan remove vlan 1 port 1
vlan add vlan 801-802 port 3
vlan remove vlan 1 port 3
vlan add vlan 801-802,901-902 port 4
vlan remove vlan 1 port 4

ring-protection logical-ring create logical-ring-name LR1 ring-id 1 west-port
4 east-port 1 WTR 1
ring-protection virtual-ring create virtual-ring-name VR1 logical-ring LR1
raps-vid 800
ring-protection virtual-ring add ring VR1 vid 801-802
ring-protection virtual-ring create virtual-ring-name VR2 logical-ring LR1
raps-vid 900
ring-protection virtual-ring add ring VR2 vid 901-902

```

## Procedure 23-4

### MPLS-TP configuration example

This example shows a general MPLS-TP configuration between DUTA, which has a loopback address of 1.1.1.1, and DUTB, which has a loopback address of 2.2.2.2.

All IP interfaces are enabled for RSVP-TE and RSVP-TE and LDP are globally enabled.

**Note 1:** The 5410 does not support GMPLS configuration.

**Note 2:** This procedure assumes that the device has already been configured with base LSR functionality as illustrated in [“Configuring a 39XX/51XX LSR” on page 14-2](#).

#### DUTA configuration

This example shows the configuration on DUTA.

```
gmpls tp-tunnel create rsvp-ingress-unidir DUTA-DUTB-1 dest-ip 2.2.2.2
gmpls tp-tunnel create static-ingress-assoc DUTA-DUTB-ASSOC-1 forward-tunnel
DUTA-DUTB-1 reverse-dyntun-name DUTB-DUTA-1

mpls l2-vpn create dynamic-vc DUTA-DUTB-VC-1 peer 2.2.2.2 pw-id 12 tp-tunnel-
assoc DUTA-DUTB-ASSOC-1
mpls l2-vpn create dynamic-vc DUTA-DUTB-VC-2 peer 2.2.2.2 pw-id 1212 tp-
tunnel-assoc DUTA-DUTB-ASSOC-1

virtual-switch Ethernet create vs VPLS1 mode vpls
virtual-switch Ethernet create vs VPWS1 mode vpws

virtual-switch Ethernet attach mpls-vc DUTA-DUTB-VC-1 vs VPLS1
virtual-switch Ethernet attach mpls-vc DUTA-DUTB-VC-2 vs VPWS1

virtual-switch Ethernet add vs VPLS1 port 1 vlan 1001
virtual-switch Ethernet add vs VPWS1 port 1 vlan 1002
```

#### DUTB configuration

This example shows the configuration on DUTB.

```
gmpls tp-tunnel create rsvp-ingress-unidir DUTB-DUTA-1 dest-ip 1.1.1.1
gmpls tp-tunnel create static-ingress-assoc DUTB-DUTA-ASSOC-1 forward-tunnel
DUTB-DUTA-1 reverse-dyntun-name DUTA-DUTB-1

mpls l2-vpn create dynamic-vc DUTB-DUTA-VC-1 peer 1.1.1.1 pw-id 12 tp-tunnel-
assoc DUTB-DUTA-ASSOC-1
mpls l2-vpn create dynamic-vc DUTB-DUTA-VC-2 peer 1.1.1.1 pw-id 1212 tp-
tunnel-assoc DUTB-DUTA-ASSOC-1

virtual-switch Ethernet create vs VPLS1 mode vpls
virtual-switch Ethernet create vs VPWS1 mode vpws
```

```
virtual-switch Ethernet attach mpls-vc DUTB-DUTA-VC-1 vs VPLS1
virtual-switch Ethernet attach mpls-vc DUTB-DUTA-VC-1 vs VPWS1

virtual-switch Ethernet add vs VPLS1 port 1 vlan 1001
virtual-switch Ethernet add vs VPWS1 port 1 vlan 1002
```

## Procedure 23-5

### End-to-End EPL service over MPLS example

This chapter describes a sample EPL over MPLS topology including the commands for configuring this example.

Two Service Aggregation Switches, a 5160 and a 5142, are connected using an MPLS topology.

Figure 23-4 shows the sample EPL over MPLS topology.

**Figure 23-4**  
Sample EPL service over MPLS topology



This example provides the commands to configure the sample end-to-end EPL service for the following switches:

- [5160 configuration](#)
- [5142 configuration](#)

### 5160 configuration

Configure MPLS LSP:

```

gmpls tp-tunnel create rsvp-ingress-unidir Test123 logical-id 1
dest-ip 8.8.8.8 explicit-tunnel-path PathA
gmpls tp-tunnel create static-ingress-assoc Test123 logical-id
1 forward-tunnel Test123 reverse-dyntun-name Test123 bfd-
monitor enable

```

Configure MPLS I2-vpn:

```

mpls l2-vpn create dynamic-vc Test123 pw-id 123 peer 8.8.8.8 tp-
tunnel-assoc Test123

```

Configure virtual switch:

```

virtual-switch ethernet create vs vsi1 mode vpls

```

Attach mpls-vc to virtual switch:

```

virtual-switch ethernet attach mpls-vc Test123 vs vsi1

```

Attach EPL AC to virtual switch:

```
virtual-switch ethernet add vs vsi1 port 1
```

**Note:** When an MPLS EPL is added to a UNI port, this UNI port is automatically removed from the default VLAN. For an EVPL, you can remove the UNI Port from the default VLAN by using the CLI. When the default VLAN membership for the EPL/EVPL port has been removed, traffic that is flooded on the default VLAN no longer egresses the UNI port.

Set I2-transform on EPL AC:

```
virtual-switch ethernet set port 1 vs vsi1 ingress-l2-transform
pop egress-l2-transform push-88a8.200
```

## 5142 configuration

Configure MPLS LSP:

```
gmpls tp-tunnel create rsvp-ingress-unidir Test123 logical-id 1
dest-ip 5.5.5.5 explicit-tunnel-path PathA
gmpls tp-tunnel create static-ingress-assoc Test123 logical-id
1 forward-tunnel Test123 reverse-dyntun-name Test123 bfd-
monitor enable
```

Configure MPLS I2-vpn:

```
mpls l2-vpn create dynamic-vc Test123 pw-id 123 peer 5.5.5.5 tp-
tunnel-assoc Test123
```

Configure virtual switch:

```
virtual-switch ethernet create vs vsi1 mode vpls
```

Attach mpls-vc to virtual switch

```
virtual-switch ethernet attach mpls-vc Test123 vs vsi1
```

Attach EPL AC to virtual switch:

```
virtual-switch ethernet add vs vsi1 port 1
```

**Note:** When an MPLS EPL is added to a UNI port, this UNI port is automatically removed from the default VLAN. For an EVPL, you can remove the UNI Port from the default VLAN by using the CLI. When the default VLAN membership for the EPL/EVPL port has been removed, traffic that is flooded on the default VLAN no longer egresses the UNI port.

Set I2-transform on EPL AC:



## 23-22 Sample MPLS topology

---

```
virtual-switch ethernet set port 1 vs vs1 ingress-l2-transform
pop egress-l2-transform push-88a8.200
```

## Procedure 23-6

### Resource affinity configuration for a TE interface example

The examples in this section show TE configurations for resource affinities, SRLG and SRLG diversity.

#### Resource affinity configuration for TE interface

This example shows the creation a rule-based resource affinity:

- Configure 32-bit mask for a TE interface.
- Define color and set bit-index.
- Create color-group-list and add colors to it to prepare a bit-mask.
- Add maximum of 32 colors to a particular resource color group.
- Assign this color-group-list (bit-mask) to a TE interface.

**Note:** It is not possible to change the affinity if an LSP passes through the TE link using the affinity value.

```
mpls traffic-eng create/delete resource-color <color-name> index <1..32>
mpls traffic-eng create/delete/add/remove* resource-color-group <color-group-name> colors <color-name-list>
mpls traffic-eng set/unset ip-interface <ip-interface-name> resource-color-group <color-group-name>
```

#### TE link on numbered interface

This example shows the MPLS configuration for a TE link on a numbered interface:

```
mpls traffic-eng set ip-interface <ip-interface-name>
link-cost <32-bit-val>]
```

#### SRLG diversity configuration for protected LSP

This is how SRLG diversity for an end-to-end protected LSP works:

- Primary LSP is signaled with specific TE requirements over the CSPF calculated ERO path.
- Once the primary LSP is operational, DCL creates a SRLG ARHOP table with a list of SRLGs for each link that the LSP is passing through.
- This list is fed to backup CSPF calculations to avoid links that intersect the primary SRLG list.

- The SRLG-disjoint is best effort for maximal SRLG-diversity. If strict SRLG diversity is not possible, intersecting SRLG link(s) is/are selected in order to establish a protecting path.
- Maximal SRLG diversity is an LSP path that has the least number of intersecting links with the same SRLG.
- DCL auto-backup gives precedence to SRLG based diversity. If none of the links of the primary LSP have an associated SRLG value then it uses node-diversity.

This example shows the creation of an SRLG for an end-to-end protected LSP:

```
MPLS setting TE information for (numbered) interface :
mpls traffic-eng set ip-interface <ip-interface>
srlg <32-bit-val1, 32-bit-val2, ..., 32-bit-valn>] ; up to 40
```

## Resource affinity configuration created using a 32-bit mask value example

This example configuration shows the affinity preferences created by a 32-bit mask value and its application for link selection:

- Configure resource affinity preferences for LSP, as well as a policy  
Include-all (the link must have colors specified)  
Link-affinity + Lsp-affinity results in Lsp-affinity.
- Include-any (the selected link must have at least one color specified)  
In this case, the Link-affinity + Lsp-affinity is TRUE.
- Exclude-any (selected link must not have any colors specified)  
In this case, the Link-affinity + Lsp-affinity is FALSE.

**Note:** Once the LSP is established, it is not possible to modify it. In addition, the auto-backup tunnel will use the same affinity and policy as the primary.

```
gmpls tp-tunnel create rsvp-ingress-corout <tunnel-name> dest-ip <ip-addr>
[resource-include-all <color-group-name>]
[resource-include-any <color-group-name>]
[resource-exclude-any <color-group-name>]
```

## Procedure 23-7

### Encapsulation configuration examples

This procedure includes a few encapsulation configuration scenarios.

#### Sample topology

This topology was used for most of the TDM testing:

**Figure 23-5**  
TDM testing topology




In most of the TDM testing, a two-DUT topology is used in which EXFO (TDM TGEN) is connected to the TDM DUT port to create a back-to-back topology.

A TDM DUT port can be an SFP port or a 3926m E1/T1 TDM FRU port.

This four-DUT topology was used for MPLS:


**Figure 23-6**  
MPLS testing topology



#### Configuration steps

These drawings illustrate the configuration steps needed to configure TDM SFP and FRU ports, respectively:

**Figure 23-7**  
**Configuring TDM SFP and FRU ports**



## Payload testing on TDM FRU

A TDM profile needs to be created and attached to the TDM VC for TDM FRU ports. The TDM profile has payload-size as a mandatory parameter, which is set using the following formula:

- For SAToP Payload = 256(E1), 192(T1)
- For CESoP Payload = 8 \* N (no. of timeslots) \* Delay

Based on the formula, these possible delay (in ms) values are supported and verified:

- 0.125

- 0.25
- 0.375
- 0.5
- 0.625
- 0.75
- 0.875
- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8

**Note:** The delay must be from 0.125 ms to 1 ms in 0.125 ms increments or from 1 ms to 8 ms in 1 ms increments. Note that this is incompatible with the 3932 platform, which supports 0.125 ms increments up to 8 ms.

## Procedures

Procedures for configuring PWE module ports are:

- [“TDM MEF8 configurations with SAToP encapsulated services over cross-connect” on page 23-27](#)
- [“TDM MEF8 configurations with CESoP encapsulated services over MPLS for dynamic co-routed tunnels” on page 23-28](#)
- [“Native MPLS configurations with CESoP encapsulated services on static tunnel and static MPLS PW” on page 23-29](#)
- [“Dry Martini services for FRU and PWE module ports” on page 23-29](#)

## TDM MEF8 configurations with SAToP encapsulated services over cross-connect

```
port tdm set mode etsi
```

**Note 1:** Set the TDM mode to either ETSI or ANSI.

```
vlan create vlan 200
vlan add vlan 200 port 2
port enable port 2
virtual-circuit ethernet create vc vcEth vlan 200
```

**Note 2:** PW configuration to make TDM traffic go through Ethernet.

```
lldp disable
rstp disable port 7
mstp disable port 7
```

**Note 3:** Disable all Ethernet protocols on the port, which must be a TDM port.

```
port tdm oc3-stm1 create oc3-stm1 poTdm1 port 7 mac port-mac
```

**Note 4:** Create the TDM SFP port.

```
port tdm enable port poTdm1
attachment-circuit satop create ac acTdm1 port poTdm1 channel 0 channel-type
dsl-e1
```

**Note 5:** SAToP encapsulated services are used here to convert TDM to packet.

```
virtual-circuit tdm-mef8 create vc vcTdm1 ac acTdm1 peer-mac 00:19:3a:00:5b:a4
in-ecn 100 out-ecn 100
```

**Note 6:** The TDM PW on which the traffic goes.

**Note 7:** The peer MAC should be the TDM port source MAC of the peer DUT.

```
virtual-switch cross-connect create xc xcTdm1 tdm-vc vcTdm1 eth-vc vcEth
```

**Note 8:** Create the cross-connect to forward the TDM PW traffic to the Ethernet.

## TDM MEF8 configurations with CESoP encapsulated services over MPLS for dynamic co-routed tunnels

A bidirectional tunnel and a unidirectional tunnel are configured in this example.

### Bidirectional tunnel

```
port tdm set mode etsi
vlan create vlan 200-201
interface create loopback LBK ip 3.3.3.3
interface create ip-interface int1 ip 10.10.10.1/30 ip-forwarding on vlan 200
interface create ip-interface int2 ip 11.11.11.1/30 ip-forwarding on vlan 201
virtual-circuit tdm create tdm-profile tdmProfile pwType cesop payload-size 7
vlan add vlan 200-201 port 1.2
rsvp-te enable ip-interface LBK
rsvp-te enable ip-interface int1
rsvp-te enable ip-interface int2
rsvp-te enable
isis instance create isis-instance ISIS level L1 area 49.0001
isis interface attach ip-interface int1 isis-instance ISIS level L1
isis interface attach ip-interface int2 isis-instance ISIS level L1
gmpls tp-tunnel create rsvp-ingress-corout cor1 dest-ip 4.4.4.4 auto-backup on
mpls l2-vpn create static-vc mplsVc pw-id 123 peer 4.4.4.4 in-label 5000 out-
label 5000 tp-tunnel-ingr-corout cor1 pw-type tdm tdm-profile tdmProfile
port tdm e1 set port tdm01 framing basic
port tdm enable port tdm01
```

```
attachment-circuit cesop create ac acTdm1 port tdm01 channel 1
virtual-circuit tdm-mpls create vc vcTdm1 ac acTdm1 tdm-profile tdmProfile
virtual-switch cross-connect create xc xcTdm1 tdm-vc vcTdm1 mpls-vc mplsVc
```

### Unidirectional tunnel

```
vlan create vlan 200
interface create loopback LBK ip 3.3.3.3
interface create ip-interface int1 ip 10.10.10.1/30 ip-forwarding on vlan 200
port tdm set mode etsi
virtual-circuit tdm create tdm-profile tdmProfile pwType cesop payload-size 128
vlan add vlan 200 port 1.2
virtual-switch ethernet create vs VS-vpls mode vpls
rsvp-te enable ip-interface LBK
rsvp-te enable ip-interface int1
gmpls tp-tunnel create rsvp-ingress-unidir cor1 dest-ip 4.4.4.4
gmpls tp-tunnel create static-ingress-assoc corAssoc forward-tunnel cor1
reverse-dyntun-name cor2 bfd-monitor enable
rsvp-te enable
mpls l2-vpn create static-vc mplsVc pw-id 123 peer 4.4.4.4 in-label 5000 out-label 6000 tp-tunnel-assoc corAssoc
isis instance create isis-instance ISIS level L1 area 49.0001
isis interface attach ip-interface int1 isis-instance ISIS level L1
aggregation disable
virtual-switch ethernet attach mpls-vc mplsVc vs VS-vpls
port tdm e1 set port tdm01 framing basic
port tdm enable port tdm01
attachment-circuit cesop create ac acTdm1 port tdm01 channel 1-2
virtual-circuit tdm-mef8 create vc vcTdm ac acTdm1 tdm-profile tdmProfile
peer-mac 00:23:8a:fa:9a:0a in-ecn 100 out-ecn 101
virtual-switch ethernet add vs VS-vpls tdm-vc vcTdm vlan 1000
```

### Native MPLS configurations with CESoP encapsulated services on static tunnel and static MPLS PW

```
port tdm set mode etsi
vlan create vlan 200
interface create loopback LBK ip 3.3.3.3
interface create ip-interface int1 ip 10.10.10.1/30 ip-forwarding on vlan 200
virtual-circuit tdm create tdm-profile tdmProfile pwType cesop payload-size 7
vlan add vlan 200 port 1.2
gmpls tp-tunnel create static-ingress-corout cor1 dest-ip 4.4.4.4 next-hop-ip 10.10.10.2 forward-out-label 200 reverse-in-label 201
mpls l2-vpn create static-vc mplsVc pw-id 200 peer 4.4.4.4 in-label 5000 out-label 5000 tp-tunnel-ingr-corout cor1 pw-type tdm tdm-profile tdmProfile
port tdm e1 set port tdm01 framing basic
port tdm enable port tdm01
attachment-circuit cesop create ac acTdm1 port tdm01 channel 1
virtual-circuit tdm-mpls create vc vcTdm1 ac acTdm1 tdm-profile tdmProfile
virtual-switch cross-connect create xc xcTdm1 tdm-vc vcTdm1 mpls-vc mplsVc
```

### Dry Martini services for FRU and PWE module ports

```
port tdm set mode ansi
vlan create vlan 200
```



### 23-30 Sample MPLS topology

---

```
virtual-circuit tdm create tdm-profile tdmProfile pwType cesop payload-size 240
```

**Note 1:** The profile is compulsory for FRU ports.

```
vlan add vlan 200 port 1.1
virtual-circuit ethernet create vc vcEth vlan 200
port tdm t1 set port tdm01 framing extended-super-frame
```

**Note 2:** The user needs to set framing in case of CESoP encapsulation.

```
port tdm enable port tdm01
```

**Note 3:** There is no need to create TDM ports because they already exist.

```
attachment-circuit cesop create ac acTdm1 port tdm01 channel 10-20
virtual-circuit tdm-dry-martini create vc vcTdm1 ac acTdm1 tdm-profile
tdmProfile peer-mac 00:23:8a:fa:9a:0a in-label 8000 out-label 8001
virtual-switch cross-connect create xc xcTdm1 tdm-vc vcTdm1 eth-vc vcEth
```



# 39XX/51XX Service Delivery, Aggregation and Virtualization Switches

## MPLS Configuration

Copyright© 2019 Ciena® Corporation. All rights reserved.

SAOS 6.18.1

Publication: 009-3313-041

Document status: Standard

Revision A

Document release date: June 2019

### **CONTACT CIENA**

For additional information, office locations, and phone numbers, please visit the Ciena web site at **[www.ciena.com](http://www.ciena.com)**