



# 7100-Series Configuration Guide

*Firmware Version 8.31*

## Legal Notices

Extreme Networks, Inc., on behalf of or through its wholly-owned subsidiary, Enterasys Networks, Inc., reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

## Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information about Extreme Networks trademarks, go to:  
[www.extremenetworks.com/company/legal/trademarks/](http://www.extremenetworks.com/company/legal/trademarks/)

## Support

For product support, including documentation, visit:  
<http://support.extremenetworks.com/>

## Contact

Extreme Networks, Inc.  
145 Rio Robles  
San Jose, CA 19534  
Tel: +1 408-579-2800  
Toll-free: +1 888-257-3000

## **Enterasys Networks, Inc. Firmware License Agreement**

### **BEFORE OPENING OR UTILIZING THE ENCLOSED PRODUCT, CAREFULLY READ THIS LICENSE AGREEMENT.**

This document is an agreement ("Agreement") between the end user ("You") and Enterasys Networks, Inc., a wholly-owned subsidiary of Extreme Networks, Inc., on behalf of itself and its Affiliates (as hereinafter defined) ("Enterasys"), that sets forth Your rights and obligations with respect to the Enterasys software program/firmware (including any accompanying documentation, hardware or media) (collectively, the "Program") in the package and prevails over any additional, conflicting or inconsistent terms and conditions appearing on any purchase order or other document submitted by You. "Affiliate" means any person, partnership, corporation, limited liability company, other form of enterprise that directly or indirectly through one or more intermediaries, controls, or is controlled by, or is under common control with the party specified.

This Agreement constitutes the entire understanding between the parties, with respect to the subject matter of this Agreement. The Program may be contained in firmware, chips or other media.

BY INSTALLING OR OTHERWISE USING THE PROGRAM, YOU REPRESENT THAT YOU ARE AUTHORIZED TO ACCEPT THESE TERMS ON BEHALF OF THE END USER (IF THE END USER IS AN ENTITY ON WHOSE BEHALF YOU ARE AUTHORIZED TO ACT, "YOU" AND "YOUR" SHALL BE DEEMED TO REFER TO SUCH ENTITY) AND THAT YOU AGREE THAT YOU ARE BOUND BY THE TERMS OF THIS AGREEMENT, WHICH INCLUDES, AMONG OTHER PROVISIONS, THE LICENSE, THE DISCLAIMER OF WARRANTY AND THE LIMITATION OF LIABILITY. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT OR ARE NOT AUTHORIZED TO ENTER INTO THIS AGREEMENT, ENTERASYS IS UNWILLING TO LICENSE THE PROGRAM TO YOU AND YOU AGREE TO RETURN THE UNOPENED PRODUCT TO ENTERASYS OR YOUR DEALER, IF ANY, WITHIN TEN (10) DAYS FOLLOWING THE DATE OF RECEIPT FOR A FULL REFUND.

IF YOU HAVE ANY QUESTIONS ABOUT THIS AGREEMENT, CONTACT EXTREME NETWORKS, LEGAL DEPARTMENT AT (408) 579-2800.

You and Enterasys agree as follows:

1. LICENSE. You are granted a revocable, non-exclusive and non-transferable right to use only one (1) copy of the Program provided in this package subject to the terms and conditions of this Agreement.
2. RESTRICTIONS. Except as otherwise authorized in writing by Enterasys, You may not, nor may You permit any third party to:
  - (a) Reverse engineer, decompile, disassemble or modify the Program, in whole or in part, including for reasons of error correction or interoperability, except to the extent expressly permitted by applicable law and to the extent the parties shall not be permitted by that applicable law, such rights are expressly excluded. Information necessary to achieve interoperability or correct errors is available from Enterasys upon request and payment of Enterasys' applicable fee.
  - (b) Incorporate the Program in whole or in part, in any other product or create derivative works based on the Program, in whole or in part.
  - (c) Publish, disclose, copy reproduce or transmit the Program, in whole or in part.
  - (d) Assign, sell, license, sublicense, rent, lease, encumber by way of security interest, pledge or otherwise transfer the Program, in whole or in part.
  - (e) Remove any copyright, trademark, proprietary rights, disclaimer or warning notice included on or embedded in any part of the Program.

3. APPLICABLE LAW. This Agreement shall be interpreted and governed under the laws and in the state and federal courts of the State of California without regard to its conflicts of laws provisions. You accept the personal jurisdiction and venue of the Superior Court of California in Santa Clara County or the United States District Court for the Northern District of California in San Jose, California. None of the 1980 United Nations Convention on the Limitation Period in the International Sale of Goods, and the Uniform Computer Information Transactions Act shall apply to this Agreement.
4. EXPORT RESTRICTIONS. You acknowledge and agree that the Program and its accompanying materials/documentation are subject to the export control laws and regulations of the United States, including but not limited to the Export Administration Regulations (EAR), the International Traffic in Arms Regulations (ITAR), and the sanction regimes of the U.S. Department of Treasury, Office of Foreign Assets Control's Foreign Assets Control Regulations (FACR). You agree that You will comply with these laws and regulations.

You agree that You will not, without prior U.S. Government authorization, export, reexport, or transfer the Program, either directly or indirectly, to any country subject to a U.S. trade embargo or sanction (e.g. Cuba, N. Korea, Iran, Syria, Sudan) or to any resident or national of said countries, or to any person, organization, or entity on any of the restricted parties lists maintained by the U.S. Departments of State, Treasury, or Commerce. In addition, You agree that You will not export, reexport or transfer the Program to any end-user engaged in activities, or for any end-use, directly or indirectly related to the design, development, production, use, or stockpiling of weapons of mass destruction, e.g. nuclear, chemical, or biological weapons, and the missile technology to deliver them.

5. UNITED STATES GOVERNMENT RESTRICTED RIGHTS. The enclosed Program (i) was developed solely at private expense; (ii) contains "restricted computer software" submitted with restricted rights in accordance with section 52.227-19 (a) through (d) of the Commercial Computer Software-Restricted Rights Clause and its successors, and (iii) in all respects is proprietary data belonging to Enterasys, its Affiliates and/or its suppliers. For Department of Defense units, the Program is considered commercial computer software in accordance with DFARS section 227.7202-3 and its successors, and use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth herein.
6. DISCLAIMER OF WARRANTY. EXCEPT FOR THOSE WARRANTIES EXPRESSLY PROVIDED TO YOU IN WRITING BY ENTERASYS, ENTERASYS DISCLAIMS ALL WARRANTIES, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT WITH RESPECT TO THE PROGRAM. IF THE IMPLIED WARRANTIES MAY NOT BE DISCLAIMED BY APPLICABLE LAW, THEN ANY IMPLIED WARRANTIES ARE LIMITED IN DURATION TO THIRTY (30) DAYS AFTER DELIVERY OF THE PROGRAM TO YOU.
7. LIMITATION OF LIABILITY. IN NO EVENT SHALL ENTERASYS OR ITS AFFILIATES AND SUPPLIERS BE LIABLE FOR ANY DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS, PROFITS, BUSINESS INTERRUPTION, LOSS OF BUSINESS INFORMATION, SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR RELIANCE DAMAGES, OR OTHER LOSS) ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM, EVEN IF ENTERASYS HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THIS FOREGOING LIMITATION SHALL APPLY REGARDLESS OF THE CAUSE OF ACTION UNDER WHICH DAMAGES ARE SOUGHT.

NOTWITHSTANDING THE FOREGOING, THE CUMULATIVE LIABILITY OF ENTERASYS TO YOU FOR ALL CLAIMS RELATING TO THE PROGRAM, IN CONTRACT,

TORT OR OTHERWISE, SHALL NOT EXCEED THE TOTAL AMOUNT OF FEES PAID TO ENTERASYS BY YOU FOR THE RIGHTS GRANTED HEREIN.

8. **AUDIT RIGHTS.** You hereby acknowledge that the intellectual property rights associated with the Program are of critical value to Enterasys and its Affiliates, and, accordingly, You hereby agree to maintain complete books, records and accounts showing: (i) license fees due and paid, and (ii) the use, copying and deployment of the Program. You also grant to Enterasys and its authorized representatives, upon reasonable notice, the right to audit and examine during Your normal business hours, Your books, records, accounts and hardware devices upon which the Program may be deployed to verify compliance with this Agreement, including the verification of the license fees due and paid to Enterasys and the use, copying and deployment of the Program. Enterasys' right of examination shall be exercised reasonably, in good faith and in a manner calculated to not unreasonably interfere with Your business. In the event such an audit discovers any non-compliance with this Agreement, including copies of the Program made, used or deployed in breach of this Agreement, You shall promptly cease such unauthorized conduct, pay to Enterasys the appropriate license fees and be subject to any other available claim from Enterasys pursuant to applicable law. Enterasys reserves the right, to be exercised in its sole discretion and without prior notice, to terminate this Agreement, including the license, effective immediately, for failure to comply with this Agreement. Upon any such termination, You shall immediately cease all use of the Program and shall return to Enterasys the Program and all copies of the Program.
9. **OWNERSHIP.** This is a license agreement and not an agreement for sale. You acknowledge and agree that the Program constitutes trade secrets and/or copyrighted material of Enterasys and/or its suppliers. You agree to implement reasonable security measures to protect such trade secrets and copyrighted material. All right, title and interest in and to the Program shall remain with Enterasys and/or its Affiliates and suppliers. All rights not specifically granted to You shall be reserved to Enterasys.
10. **TRADEMARKS.** ENTERASYS, ENTERASYS NETWORKS, ENTERASYS SECURE NETWORKS, NETSIGHT, ENTERASYS NETSIGHT, and any logos associated therewith, are trademarks or registered trademarks of Enterasys Networks, Inc., in the United States and/or other countries. Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries. For more information on Enterasys and Extreme trademarks, see: [www.extremenetworks.com/about-extreme/trademarks.aspx](http://www.extremenetworks.com/about-extreme/trademarks.aspx).

All other product names mentioned in Program may be trademarks or registered trademarks of their respective companies.

11. **ENFORCEMENT.** You acknowledge and agree that any breach of this Agreement by You may cause Enterasys irreparable damage for which recovery of money damages would be inadequate, and that Enterasys may be entitled to seek timely injunctive relief to protect Enterasys' rights under this Agreement in addition to any and all remedies available at law.
12. **ASSIGNMENT.** You may not assign, transfer or sublicense this Agreement or any of Your rights or obligations under this Agreement, except that You may assign this Agreement to any person or entity that acquires substantially all of Your stock assets. Enterasys may assign this Agreement in its sole discretion to anyone or any entity without Your consent and without notice, including assigning this Agreement to its parent company, Extreme Networks, Inc. This Agreement shall be binding upon and inure to the benefit of the parties, their legal representatives, permitted transferees, successors and assigns as permitted by this Agreement. Any attempted assignment, transfer or sublicense in violation of the terms of this Agreement shall be void and a breach of this Agreement.

13. WAIVER. A waiver by Enterasys of a breach of any of the terms and conditions of this Agreement must be in writing and will not be construed as a waiver of any subsequent breach of such term or condition. Enterasys' failure to enforce a term upon Your breach of such term shall not be construed as a waiver of Your breach or prevent enforcement on any other occasion.
14. SEVERABILITY. In the event any provision of this Agreement is found to be invalid, illegal or unenforceable, the validity, legality and enforceability of any of the remaining provisions shall not in any way be affected or impaired thereby, and that provision shall be reformed, construed and enforced to the maximum extent permissible. Any such invalidity, illegality, or unenforceability in any jurisdiction shall not invalidate or render illegal or unenforceable such provision in any other jurisdiction.
15. TERMINATION. Enterasys may terminate this Agreement immediately upon Your breach of any of the terms and conditions of this Agreement. Upon any such termination, You shall immediately cease all use of the Program and shall return to Enterasys the Program and all copies of the Program.



---

# Contents

## About This Guide

How to Use This Guide .....	xxxi
Related Documents .....	xxxi
Conventions Used in This Guide .....	xxxi
Commonly Used Acronyms .....	xxxii
Getting Help .....	xxxii

## Chapter 1: Getting Started

Device Management Methods .....	1-1
Initial Configuration .....	1-1
Advanced Configuration Overview .....	1-2

## Chapter 2: Using the CLI

CLI Conventions .....	2-1
Getting Help with CLI Syntax .....	2-1
Using Context-Sensitive Help .....	2-1
Performing Keyword Lookups .....	2-2
Displaying Scrolling Screens .....	2-3
Abbreviating and Completing Commands .....	2-3
Using the Spacebar Auto Complete Function .....	2-4
Configuring CLI Properties .....	2-4
Example CLI Properties Configuration .....	2-4
CLI Properties Display Commands .....	2-5

## Chapter 3: Image Configuration and File Management

Configuration and Image File Management on Your System .....	3-1
Automated Deployment .....	3-2
Saving a Configuration .....	3-2
Executing a Configuration .....	3-3
Deleting a Configuration Restore-Point or File .....	3-3
Downloading a File from an FTP, TFTP, or SCP Server .....	3-4
Downloading a Firmware Image via the Serial Port .....	3-5
Uploading a Configuration File .....	3-8
Setting the Boot Firmware Image .....	3-8
Running a Configuration Script .....	3-8
Configuration and Image File Display Commands .....	3-9

## Chapter 4: High Availability Firmware Upgrade (HAU) Configuration

Using High Availability Firmware Upgrade in Your Network .....	4-1
Implementing HAU .....	4-2
High Availability Upgrade Preconditions .....	4-2
System Limitations During a High Availability Upgrade .....	4-3
HAU Configuration Overview .....	4-3
Configuring System Boot Image and Mode .....	4-3
Configuring HAU Default Mode .....	4-4
Configuring HAU Groups .....	4-4
Configuring a Delay Between HAU Group Upgrades .....	4-5
Disabling a Configured HAU .....	4-5
Forcing Early Completion of a Running HAU .....	4-5
High Availability Firmware Upgrade in a Virtual Switch Bonded System .....	4-6



Configuring HAU .....	4-6
Terms and Definitions .....	4-8

## Chapter 5: Virtual Switch Bonding (VSB) Configuration

Using Virtual Switch Bonding in Your Network .....	5-1
Implementing VSB Stacking .....	5-4
VSB Configuration Overview .....	5-4
VSB Chassis Configuration .....	5-5
VSB Interconnect Link Configuration .....	5-5
Link Failure Response (LFR) Configuration .....	5-6
VSB System MAC address Configuration .....	5-8
Licensing .....	5-9
Globally Enabling and Disabling the VSB System .....	5-9
High Availability Firmware Upgrade .....	5-9
Applying a VSB Configuration File to a Replacement VSB Chassis .....	5-10
Adding a Chassis to an Existing VSB Stacking System .....	5-10
A Failed Chassis in an VSB Stacking System .....	5-10
Configuring VSB .....	5-11
Terms and Definitions .....	5-12

## Chapter 6: Port Configuration

Port Configuration Overview .....	6-1
Port String Syntax Used in the CLI .....	6-2
Console Port Parameters .....	6-3
Administratively Enabling a Port .....	6-4
Ingress Filtering .....	6-4
Port Alias .....	6-4
Force Linkdown .....	6-4
Default Port Speed .....	6-5
Port Duplex .....	6-5
Jumbo Frames .....	6-6
Auto-Negotiation and Port Advertised Ability .....	6-7
Port MDI/MDIX .....	6-8
Port Flow Control .....	6-8
Configuring Link Traps and Link Flap Detection .....	6-9
Port Broadcast Suppression .....	6-10
Port Priority .....	6-10
Transmit Queue Monitoring .....	6-11
Energy Efficient Ethernet (EEE) .....	6-12
Configuring Ports .....	6-12
Terms and Definitions .....	6-17

## Chapter 7: Ethernet Operations, Administration, and Maintenance (OAM) Configuration

Using Ethernet OAM in Your Network .....	7-1
Implementing Ethernet OAM .....	7-2
Ethernet OAM Overview .....	7-2
OAM Client .....	7-2
OAM Discovery .....	7-2
OAM Client Mode .....	7-3
OAM Datalink Layer Monitoring .....	7-3
OAM Event Notification Retries .....	7-7
Unidirectional Link Detection (ULD) .....	7-7
Configuring Ethernet OAM .....	7-8
Ethernet OAM Configuration Example .....	7-10
Terms and Definitions .....	7-11

## Chapter 8: Port Mirroring Configuration

How to Use Port Mirroring in Your Network .....	8-1
Implementing Port Mirroring .....	8-3
Overview of Port Mirroring Configurations .....	8-3
LAG Mirrors .....	8-3
VLAN Mirrors .....	8-4
Configuring Port Mirrors .....	8-4
Reviewing Port Mirroring .....	8-4
Setting Port or VLAN Mirroring .....	8-5
Deleting Mirrors .....	8-6
Remote Mirroring Using a Layer 2 GRE Tunnel .....	8-6
Example: Configuring and Monitoring Port Mirroring .....	8-8

## Chapter 9: System Configuration

System Properties Overview .....	9-1
System Properties Example .....	9-4
User Management Overview .....	9-5
User Management Example .....	9-7
Setting the Authentication Login Method .....	9-7
Using WebView .....	9-8
Management Authentication Notification MIB Overview .....	9-8
Configuring Management Authentication Notification MIB .....	9-8
Management Authentication Notification MIB Configuration Examples .....	9-9
License Overview .....	9-10
Configuring a License .....	9-10
SNTP Overview .....	9-11
Unicast Polling Mode .....	9-11
Broadcast Listening Mode .....	9-11
SNTP Authentication .....	9-11
Configuring SNTP .....	9-13
SNTP Configuration Examples .....	9-15
Telnet Overview .....	9-16
Configuring Telnet .....	9-17
Telnet Examples .....	9-17
Secure Shell Overview .....	9-17
SSH Client Authentication .....	9-18
Configuring Secure Shell .....	9-19
Secure Shell Configuration Examples .....	9-20
Domain Name Server (DNS) Overview .....	9-21
Configuring DNS .....	9-21
DNS Configuration Example .....	9-22
DHCP Overview .....	9-23
IPv4 DHCP Supported Server Options .....	9-24
DHCP Server .....	9-26
DHCPv6 Overview .....	9-28
DHCPv6 Server Option Information Configuration Example .....	9-29
IPv6 DHCP Relay Source and Destination Interfaces .....	9-30
Configuring DHCP .....	9-31
Node Alias Overview .....	9-34
Configuring Node Alias .....	9-34
Setting Node Alias State and Max Entries .....	9-35
MAC Address Settings Overview .....	9-36
Age Time .....	9-37
Multicast MAC Address VLAN Port Limit .....	9-37
Network Load Balanced (NLB) Servers Configured for Multicast .....	9-37

Static MAC Address Entry .....	9-38
Unicast as Multicast .....	9-38
New and Moved MAC Address Detection .....	9-38
Terms and Definitions .....	9-39

## Chapter 10: Security Mode Configuration

How to Use Security Mode in Your Network .....	10-1
FIPS Security Mode .....	10-1
Security Profile Mode .....	10-2
Boot Access Security Mode .....	10-2
Security Profile Mode Default Parameter Setting Changes .....	10-3
Security Profile Mode Parameter Range Changes .....	10-4
C2 Security Profile Mode Command Access Changes .....	10-4
C2 Security Profile Mode Read-Write User Mode Changes .....	10-4
C2 Security Profile Mode Read-Only User Mode Changes .....	10-6
Implementing Security Mode .....	10-6
Configuring Security Mode .....	10-6
Security Mode Display Commands .....	10-7
Security Mode Configuration Example .....	10-7
Terms and Definitions .....	10-7

## Chapter 11: IPsec Protocol Configuration

How to Use IPsec in Your Network .....	11-1
IPsec Implementation Requirements .....	11-2
Required Manual Configuration .....	11-2
Understanding the IPsec Protocol .....	11-3
IKE Map .....	11-3
Configuring IPsec .....	11-7
IKE Proposal Configuration .....	11-7
IKE Policy Configuration .....	11-8
IKE Map Configuration .....	11-8
IPsec Configuration .....	11-9
IPsec Display Commands .....	11-10
IPsec Configuration Example .....	11-10
Terms and Definitions .....	11-12

## Chapter 12: Tracked Object Manager Configuration

Using Tracked Object Manager in Your Network .....	12-1
Tracked Objects .....	12-1
Probes .....	12-2
Scheduling .....	12-2
State Probe Configuration .....	12-3
Probe Parameters .....	12-4
Fail Detection Methods .....	12-5
Preset Default ICMP Probes .....	12-6
Configuring a Probe for Policy Based Routing .....	12-7
Configuring a Probe for VRRP .....	12-8
Configuring State Probes .....	12-8
Timing Probe Configuration .....	12-10
Timing Probe Parameters .....	12-10
Configuring a Timing Probe for IP SLA .....	12-11
Procedure .....	12-12
Tracked Object Configuration .....	12-12
Tracked Object Parameters .....	12-12
Procedure .....	12-13

Example .....	12-14
Terms and Definitions .....	12-15

## Chapter 13: IP SLA Configuration

Using IP SLA in Your Network .....	13-1
Constraints and Limitations .....	13-1
Monitoring Paths .....	13-2
Scheduling Tests .....	13-2
Reported Statistics .....	13-2
Measurements .....	13-3
System Resources Affected by IP SLA .....	13-3
IP SLA Syslog Messages .....	13-4
Configuring IP SLA .....	13-4
Default Settings .....	13-5
IP SLA Configuration Procedure .....	13-6
Example IP SLA Configuration .....	13-6
IP SLA Display Commands .....	13-7

## Chapter 14: Power over Ethernet Configuration

How to Use PoE in Your Network .....	14-1
Implementing PoE .....	14-1
Allocation of PoE Power to Modules .....	14-2
Management of PoE Power to PDs .....	14-3
Configuring PoE .....	14-3
Default Settings .....	14-3
PoE Configuration Procedure .....	14-4
PoE Display Commands .....	14-6

## Chapter 15: Discovery Protocol Configuration

How to Use Neighbor Discovery in Your Network .....	15-1
Understanding Neighbor Discovery .....	15-2
LLDP-MED .....	15-3
LLDPDU Frames .....	15-5
Neighbor Warning Detection .....	15-7
Configuring LLDP .....	15-7
LLDP Configuration Commands .....	15-7
Basic LLDP Configuration .....	15-9
LLDP Display Commands .....	15-10
Configuring Neighbor Warning Detection .....	15-12
Configuring Enterasys Discovery Protocol .....	15-12
Enterasys Discovery Protocol Configuration Commands .....	15-12
Enterasys Discovery Protocol Show Commands .....	15-13
Configuring Cisco Discovery Protocol .....	15-13
Cisco Discovery Protocol Configuration Commands .....	15-13
Cisco Discovery Protocol Show Commands .....	15-13

## Chapter 16: Data Center Bridging Configuration

How to Use Data Center Bridging in Your Network .....	16-1
Implementing Data Center Bridging .....	16-2
Enhanced Transmission Selection Configuration .....	16-2
Priority-Based Flow Control Configuration .....	16-4
Application Priority Configuration .....	16-5
Congestion Notification (CN) Configuration .....	16-5
Implementing Congestion Notification .....	16-7

Enabling Congestion Notification .....	16-8
Congestion Notification Priority Value (CNPV) .....	16-8
Congestion Notification Domain Defense .....	16-9
LLDP .....	16-12
Congestion Point Queue .....	16-13
Congestion Notification Queue Profile .....	16-13
Congestion Notification Configuration Example .....	16-14
Configuring Data Center Bridging .....	16-14
Terms and Definitions .....	16-17

## Chapter 17: Simple Network Management Protocol (SNMP) Configuration

Using SNMP in Your Network .....	17-1
High-Level Configuration Process .....	17-2
SNMP Concepts .....	17-2
Manager/Agent Model Components .....	17-2
Message Functions .....	17-2
Access to MIB Objects .....	17-3
SNMP Support on 7100-Series Devices .....	17-4
Versions Supported .....	17-4
Terms and Definitions .....	17-5
Security Models and Levels .....	17-6
Access Control .....	17-7
Configuring SNMP .....	17-7
Configuration Basics .....	17-8
How SNMP Processes a Notification Configuration .....	17-8
SNMP Defaults .....	17-9
Configuring SNMPv1/SNMPv2c .....	17-9
Configuring SNMPv3 .....	17-11
Configuring Secure SNMP Community Names .....	17-18
Reviewing SNMP Settings .....	17-20
Community .....	17-20
Context .....	17-21
Counters .....	17-21
Engineid .....	17-22
Groups .....	17-22
Group Access Rights .....	17-23
Target Parameter Profiles .....	17-24
Target Address Profiles .....	17-24
Notify .....	17-24
Notify Filter .....	17-25
Notify Profile .....	17-25
Users .....	17-26
Views .....	17-26

## Chapter 18: Spanning Tree Configuration

What Is the Spanning Tree Protocol? .....	18-1
Why Would I Use Spanning Trees in My Network? .....	18-2
How Do I Implement Spanning Trees? .....	18-3
STP Overview .....	18-3
Rapid Spanning Tree .....	18-4
Multiple Spanning Tree .....	18-4
Functions and Features Supported on the 7100-Series Device .....	18-6
Spanning Tree Versions .....	18-6
Maximum SID Capacities .....	18-6
Network Diameter .....	18-6

Port Forwarding .....	18-6
Disabling Spanning Tree .....	18-7
STP Features .....	18-7
Multisource Detection .....	18-9
Understanding How Spanning Tree Operates .....	18-9
Spanning Tree Basics .....	18-10
Electing the Root Bridge .....	18-10
Assigning Path Costs .....	18-10
Paths to Root .....	18-10
Identifying Designated, Alternate, and Backup Port Roles .....	18-12
Assigning Port States .....	18-14
RSTP Operation .....	18-14
MSTP Operation .....	18-15
Multisource Detection .....	18-20
Configuring STP and RSTP .....	18-21
Reviewing and Enabling Spanning Tree .....	18-21
Adjusting Spanning Tree Parameters .....	18-22
Enabling the Backup Root Function .....	18-25
Adjusting RSTP Parameters .....	18-26
Configuring MSTP .....	18-27
Example 1: Configuring MSTP for Traffic Segregation .....	18-28
Example 2: Configuring MSTP for Maximum Bandwidth Utilization .....	18-30
Adjusting MSTP Parameters .....	18-31
Monitoring MSTP .....	18-32
Understanding and Configuring SpanGuard .....	18-32
What Is SpanGuard? .....	18-32
How Does It Operate? .....	18-33
Configuring SpanGuard .....	18-33
Understanding and Configuring Loop Protect .....	18-34
What Is Loop Protect? .....	18-34
How Does It Operate? .....	18-34
Configuring Loop Protect .....	18-37
Terms and Definitions .....	18-39

## Chapter 19: Shortest Path Bridging (SPB) Configuration

Using Shortest Path Bridging (SPB) in Your Network .....	19-1
Implementing Shortest Path Bridging .....	19-3
Shortest Path Bridging VLAN Configuration Overview .....	19-4
SPBV Spanning Tree Configuration .....	19-4
SPVID Pool .....	19-4
Assigning a Base-VLAN to Use SPB .....	19-5
Base-VLAN Configuration .....	19-5
SPB Ports .....	19-5
Configuring Shortest Path Bridging VLAN .....	19-6
Terms and Definitions .....	19-7

## Chapter 20: VLAN Configuration

Using VLANs in Your Network .....	20-1
Implementing VLANs .....	20-2
Preparing for VLAN Configuration .....	20-3
Understanding How VLANs Operate .....	20-3
Learning Modes and Filtering Databases .....	20-3
VLAN Assignment and Forwarding .....	20-4
Example of a VLAN Switch in Operation .....	20-6
VLAN Support on Extreme Networks 7100-Series Switches .....	20-6

Maximum Active VLANs .....	20-7
Configurable Range .....	20-7
VLAN Types .....	20-7
Dynamic VLAN Support .....	20-8
Configuring VLANs .....	20-9
Default Settings .....	20-10
Configuring Static VLANs .....	20-11
Creating a Secure Management VLAN .....	20-13
Configuring Dynamic VLANs .....	20-14
Configuring Protocol-Based VLAN Classification .....	20-15
Configuring IGMP VLAN Snooping .....	20-17
Monitoring VLANs .....	20-18
Terms and Definitions .....	20-18
VLAN Provider Bridges .....	20-19
Configuring Provider Bridges .....	20-22

## Chapter 21: Link Aggregation Control Protocol (LACP) Configuration

Using Link Aggregation in Your Network .....	21-1
Implementing Link Aggregation .....	21-2
Link Aggregation Overview .....	21-3
LACP Operation .....	21-3
How a LAG Forms .....	21-3
Attached Ports .....	21-5
LAG Port Parameters .....	21-7
The Out-Port Algorithm .....	21-8
Static Port Assignment .....	21-8
Platform LAG and Physical Port Support .....	21-8
Configuring Link Aggregation .....	21-9
Link Aggregation Configuration Examples .....	21-11
Link Aggregation Configuration Example 1 .....	21-11
Link Aggregation Configuration Example 2 .....	21-15
Terms and Definitions .....	21-19

## Chapter 22: Policy Configuration

Using Policy in Your Network .....	22-1
Implementing Policy .....	22-2
Policy Overview .....	22-2
Introduction .....	22-2
Understanding Roles in a Secure Network .....	22-3
Policy Roles .....	22-3
VLAN-to-Policy Mapping .....	22-5
Applying Policy Using the RADIUS Response Attributes .....	22-5
Classification Rules .....	22-7
Policy Capabilities .....	22-10
Configuring Policy .....	22-10
Policy Configuration Example .....	22-15
Roles .....	22-16
Policy Domains .....	22-17
Platform Configuration .....	22-18
Terms and Definitions .....	22-25

## Chapter 23: Multicast Configuration

How to Use Multicast in Your Network .....	23-1
Implementing Multicast .....	23-2
Understanding Multicast .....	23-2
Internet Group Management Protocol (IGMP) .....	23-2
Distance Vector Multicast Routing Protocol (DVMRP) .....	23-5
Protocol Independent Multicast (PIM) .....	23-11
Configuring Multicast .....	23-20
Configuring IGMP .....	23-20
Configuring DVMRP .....	23-23
Configuring PIM .....	23-25

## Chapter 24: Multicast Listener Discovery (MLD) Configuration

Using MLD in Your Network .....	24-1
Implementing MLD .....	24-1
Understanding MLD .....	24-2
Configuring MLD .....	24-5

## Chapter 25: System Logging Configuration

Using Syslog in Your Network .....	25-1
Syslog On 7100-Series Switches .....	25-2
Syslog Overview .....	25-2
Configuring Syslog Message Disposition .....	25-2
Filtering by Severity and Facility .....	25-3
Syslog Components and Their Use .....	25-3
Basic Syslog Scenario .....	25-4
Interpreting Messages .....	25-6
Configuring Syslog .....	25-6
Syslog Command Precedence .....	25-7
About Server and Application Severity Levels .....	25-7
Configuring Syslog Server(s) .....	25-7
Modifying Syslog Server Defaults .....	25-8
Reviewing and Configuring Logging for Applications .....	25-9
Enabling Console Logging and File Storage .....	25-10
CLI and SNMP Audit Logging .....	25-11
Configuration Examples .....	25-12

## Chapter 26: Network Monitoring Configuration

Using Network Monitoring in Your Network .....	26-1
Network Monitoring Overview .....	26-2
Console/Telnet History Buffer .....	26-2
Network Diagnostics .....	26-2
Switch Connection Statistics .....	26-3
Users .....	26-4
RMON .....	26-4
SMON Priority and VLAN Statistics Counting .....	26-5
Configuring Network Monitoring .....	26-7

## Chapter 27: Connectivity Fault Management Configuration

How to Use Connectivity Fault Management in Your Network .....	27-1
Connectivity Fault Management Overview .....	27-2
Maintenance Domain (MD) .....	27-3
Maintenance Association (MA) .....	27-4
Maintenance Point (MP) .....	27-5



CFM Configuration Modes .....	27-8
Implementing Connectivity Fault Management .....	27-8
Configuring CFM at the Global System Level .....	27-9
CFM Logging Filtering .....	27-9
VLAN Table Configuration .....	27-9
Activating CFM Configuration .....	27-11
Configuring a Maintenance Domain (MD) .....	27-11
MD Configuration Modes .....	27-11
MD Naming Conventions .....	27-13
Setting SenderID TLV Permission .....	27-13
Enabling Maintenance Intermediate-Points (MIP) .....	27-14
Setting the MD Level .....	27-14
Changing the Maintenance Domain Name .....	27-15
Configuring a Maintenance Association (MA) .....	27-15
Accessing MA Configuration Mode .....	27-16
Enabling the Maintenance Association Configuration .....	27-16
Changing the Maintenance Association Name .....	27-16
Setting the Continuity Check Message (CCM) Interval .....	27-17
Configuring the Maintenance Association MEP List .....	27-17
Configuring the Maintenance Association Components .....	27-18
Configuring a Maintenance End-Point (MEP) .....	27-18
Accessing MEP Configuration Mode .....	27-19
Configuring the MEP Bridge Port .....	27-19
Configuring the MEP VLAN .....	27-19
Configuring MEP Direction .....	27-20
Setting the Lowest Priority MEP Defect Alarm .....	27-20
Enabling MEP CCMs .....	27-21
Activating the MEP State Machine and the Remote MEP .....	27-22
Modifying the MEP CCM and Linktrace 802.1p Priority .....	27-22
Enabling the Maintenance End-point Configuration .....	27-23
CFM Loopback and Linktrace Protocols .....	27-23
The CFM Loopback Protocol .....	27-23
The CFM Linktrace Protocol .....	27-23
Configuring Connectivity Fault Management .....	27-24
Single MD Configuration Example .....	27-31
Configuring Device maCE1:1 .....	27-32
Configuring Device maCE1:2 .....	27-33
Configuring Device maCE1:3 .....	27-34
Configuring Switch 1 .....	27-35
Configuring Device maCE2:1 .....	27-36
Configuring Device maCE2:2 .....	27-37
Configuring Device maCE2:3 .....	27-38
Configuring Switch 2 .....	27-39
Multiple MD Configuration Example .....	27-39
Configuring CE Device 1 .....	27-40
Configuring CE Device 2 .....	27-42
Configuring CE Device 3 .....	27-43
Terms and Definitions .....	27-44

## Chapter 28: Virtual Routing and Forwarding (VRF) Configuration

Using VRF in Your Network .....	28-1
Implementing VRF .....	28-1
VRF Overview .....	28-1
VRFs, Interfaces, and IP Addresses .....	28-3
Configuring VRF .....	28-4

Terms and Definitions .....	28-4
-----------------------------	------

## Chapter 29: IP Routing Configuration

The Router .....	29-1
Entering Router Configuration .....	29-1
Display Router Configuration .....	29-2
The Routing Interface .....	29-3
IP Routing Addresses .....	29-4
Secondary and Private VLAN .....	29-7
Non-Forwarding IP Management Interfaces .....	29-9
Show Interface Examples .....	29-10
IP Static Routes .....	29-12
Traffic Non-Forwarding IP Static Routes .....	29-12
IPv6 Neighbor Discovery .....	29-13
Duplicate Address Detection .....	29-13
IPv6 Address Autoconfiguration .....	29-14
Binding an IPv6 Address to a MAC Hardware Address .....	29-14
Configuring IPv6 Neighbor Discovery .....	29-14
The ARP Table .....	29-15
Gratuitous ARP .....	29-15
Proxy ARP .....	29-16
ARP/ND Proxy-All .....	29-16
Removing the Multicast ARP Restriction .....	29-16
ARP Configuration Examples .....	29-16
Router Management and Information Display .....	29-18
Terms and Definitions .....	29-20

## Chapter 30: Routing Information Protocol (RIP) Configuration

Using RIP in Your Network .....	30-1
RIP Overview .....	30-1
Configuring RIP Authentication .....	30-2
Configuring RIP Offset .....	30-4
Configuring RIP .....	30-4
Terms and Definitions .....	30-5

## Chapter 31: Routing Information Protocol Next Generation (RIPng) Configuration

Using RIPng in Your Network .....	31-1
RIPng Configuration Overview .....	31-2
Configuring RIPng .....	31-3
Terms and Definitions .....	31-4

## Chapter 32: Open Shortest Path First (OSPFv2) Configuration

Using the OSPF Protocol in Your Network .....	32-1
Implementing OSPF .....	32-2
OSPF Overview .....	32-3
Configuring Basic OSPF Parameters .....	32-3
Configuring the Router ID .....	32-5
Configuring the Designated Router .....	32-6
Configuring the Administrative Distance for OSPF Routes .....	32-8
Configuring OSPF Areas .....	32-9
Configuring Route Redistribution .....	32-16
Filtering Routes from the OSPF Route Table .....	32-16
Configuring Passive Interfaces .....	32-16
Graceful Restart .....	32-17

Configuring Interface Cost .....	32-18
Configuring OSPF with Authentication at the Interface .....	32-19
Configuring Bidirectional Forwarding Detection (BFD) on Interfaces .....	32-19
Configuring OSPF Timers .....	32-20
Configuring the PE-CE Protocol .....	32-20
Configuring OSPF .....	32-22
Default Settings .....	32-22

## Chapter 33: Open Shortest Path First Version 3 (OSPFv3) Configuration

Using the OSPFv3 Protocol in Your Network .....	33-1
OSPFv3 and OSPFv2 Differences .....	33-2
OSPFv3 and OSPFv2 Similarities .....	33-4
IPsec for OSPFv3 .....	33-4
Implementing OSPFv3 .....	33-5
OSPFv3 Configuration Overview .....	33-5
Configuring Basic OSPFv3 Parameters .....	33-5
Configuring the Router ID .....	33-8
Configuring the Designated Router .....	33-8
Configuring the Administrative Distance for OSPF Routes .....	33-11
Configuring OSPFv3 Areas .....	33-11
Configuring IPsec Authentication for OSPFv3 .....	33-19
Configuring Route Redistribution .....	33-19
Filtering Routes from the OSPF Route Table .....	33-20
Configuring Passive Interfaces .....	33-20
Graceful Restart .....	33-20
Configuring Interface Cost .....	33-22
Configuring Bidirectional Forwarding Detection (BFD) on Interfaces .....	33-23
Configuring OSPFv3 Timers .....	33-23
Configuring the PE-CE Protocol .....	33-24
OSPFv3 Configuration Details .....	33-25
Default Settings .....	33-25

## Chapter 34: Intermediate System To Intermediate System (IS-IS) Configuration

Using IS-IS in Your Network .....	34-1
Implementing IS-IS .....	34-4
IS-IS Configuration Overview .....	34-4
Enabling IS-IS Globally .....	34-5
Enabling IS-IS on the Interface .....	34-5
Configuring a Network Entity Title (NET) .....	34-6
Configuring Administrative Distance .....	34-8
Configuring IS-IS Authentication .....	34-8
Configuring Multiple Parallel Routes .....	34-9
Enabling Route Summarization .....	34-9
Configuring Route Redistribution .....	34-11
Configuring IS-IS Timers .....	34-11
Configuring the TLV Metric Style .....	34-12
Configuring IS-IS Priority .....	34-12
Configuring the IS-IS Intermediate System as Overloaded .....	34-13
Configuring the IPv6 Unicast Address Family .....	34-13
Graceful Restart .....	34-13
Configuring IS-IS .....	34-14
Terms and Definitions .....	34-18

## Chapter 35: Border Gateway Protocol (BGP) Configuration

Using BGP in Your Network .....	35-1
Implementing BGP .....	35-4
BGP Overview .....	35-5
Injecting Routes Into BGP .....	35-5
Using AS-Path Regular Expressions .....	35-7
Route Selection Preference .....	35-8
Multi-Exit Discriminator (MED) .....	35-8
Route Aggregation .....	35-9
Source IP Address Update to the Peer .....	35-10
Scalability and the Peer Full Mesh Requirement .....	35-11
Outbound Route Filtering (ORF) .....	35-13
Conditional Advertisement .....	35-13
BGP Soft Reset .....	35-14
Community and Extended Community Attributes .....	35-15
Route Flap Dampening .....	35-18
Graceful Restart .....	35-21
Configuring BGP .....	35-22
Configuring Basic BGP Router Parameters .....	35-25
Configuring BGP Route Injection .....	35-26
Configuring External BGP Basic Peering .....	35-27
Configuring Internal BGP Basic Peering .....	35-29
Configuring Multihop EBGP Basic Peering .....	35-31
Configuring BGP Neighbor Parameters .....	35-34
Configuring Source IP Address Update .....	35-35
Configuring BGP Confederations .....	35-37
Configuring Route Reflection .....	35-40
Configuring Outbound Route Filtering (ORF) .....	35-43
Configuring Conditional Advertisement .....	35-43
Configuring BGP Soft Reset .....	35-46
Configuring Flap Dampening .....	35-46
Configuring Graceful Restart .....	35-50
BGP Monitoring and Clearing .....	35-51
Terms and Definitions .....	35-51

## Chapter 36: Virtual Router Redundancy Protocol (VRRP) Configuration

Using VRRP in Your Network .....	36-1
Implementing VRRP in Your Network .....	36-2
VRRP Overview .....	36-2
Basic VRRP Topology .....	36-2
VRRP Virtual Router Creation .....	36-3
VRRP Master Election .....	36-3
Configuring a VRRP Critical-IP Address .....	36-3
Configuring VRRP Authentication .....	36-5
Enabling Master Preemption .....	36-5
Enabling Fabric Route Mode on the VRRP Backup Router .....	36-5
Enabling the VRRP Virtual Router .....	36-7
Configuring VRRP .....	36-7
VRRP Configuration Examples .....	36-9
Basic VRRP Configuration Example .....	36-9
Multiple Backup VRRP Configuration Example .....	36-11
Terms and Definitions .....	36-13

## Chapter 37: Security Configuration

Using Security Features in Your Network .....	37-1
MAC Locking .....	37-1
Secure Shell .....	37-1
TACACS+ .....	37-1
Host Denial of Service (DoS) .....	37-2
Implementing Security .....	37-2
Security Overview .....	37-3
MAC Locking .....	37-3
Secure Shell .....	37-4
TACACS+ .....	37-5
Host DoS .....	37-7
Configuring Security .....	37-8
Configuring MAC Locking .....	37-8
Configuring Secure Shell .....	37-10
Configuring TACACS+ .....	37-11
Configuring Host DoS .....	37-12

## Chapter 38: Access Control List Configuration

Using Access Control Lists (ACLs) in Your Network .....	38-1
Implementing ACLs .....	38-1
ACL Overview .....	38-2
Creating an ACL .....	38-2
Creating ACL Rules .....	38-3
Managing ACL Rules .....	38-5
Applying ACLs .....	38-7
Permit Layer 4 Fragmented Packet Behavior .....	38-8
Configuring ACLs .....	38-8
Terms and Definitions .....	38-14

## Chapter 39: Route-Map Manager Configuration

Using Route-Map Manager in Your Network .....	39-1
Implementing Route-Maps .....	39-3
Implementing a Policy Based Route-Map .....	39-3
Implementing a Redistribution Route-Map .....	39-3
Implementing an OSPF Filter Route-Map .....	39-4
Implementing a BGP Route-Map .....	39-4
Route-Map Manager Overview .....	39-4
Creating a Route-Map .....	39-4
Configuring Match and Set Clauses .....	39-5
Assigning a Policy Route-Map to an Interface .....	39-7
Configuring Route-Map Manager .....	39-8
Route-Map Manager Configuration Examples .....	39-13
Policy Based Route-Map Example .....	39-13
Redistribution Route-Map Example .....	39-14
BGP Route-Map Example .....	39-15
Terms and Definitions .....	39-15

## Chapter 40: Quality of Service (QoS) Configuration

Using Quality of Service in Your Network .....	40-1
Implementing Quality of Service .....	40-2
Quality of Service Overview .....	40-2
Class of Service (CoS) .....	40-2
CoS Priority and ToS Rewrite .....	40-3

Preferential Queue Treatment for Packet Forwarding .....	40-4
Rate Limiting .....	40-6
Rate Shaping .....	40-7
Understanding QoS Configuration on the 7100-Series .....	40-8
Determining CoS Port-Type .....	40-9
Configuring CoS Port Groups .....	40-10
Configuring CoS Port-Resource .....	40-12
Configuring CoS Reference Mapping .....	40-14
Configuring the CoS Index .....	40-15
Enabling CoS State .....	40-17
The QoS CLI Command Flow .....	40-17
Terms and Definitions .....	40-18

## Chapter 41: Anti-Spoofing Configuration

Anti-Spoofing Feature Overview .....	41-1
DHCP Snooping .....	41-2
Dynamic ARP Inspection (DAI) .....	41-2
IP Source Guard .....	41-3
Duplicate IP Address Detection .....	41-3
Populating the MAC-to-IP Binding Table .....	41-3
Implementing Anti-Spoofing in Your Network .....	41-4
Using DHCP Snooping Only .....	41-4
Anti-Spoofing Configuration .....	41-5
Overview .....	41-5
Configuration Examples .....	41-6

## Chapter 42: RADIUS-Snooping Configuration

Using RADIUS-Snooping in Your Network .....	42-1
Implementing RADIUS-Snooping .....	42-2
RADIUS-Snooping Overview .....	42-2
RADIUS-Snooping Configuration .....	42-2
RADIUS-Snooping Management .....	42-3
RADIUS Session Attributes .....	42-4
Configuring RADIUS-Snooping .....	42-5
Configuring RADIUS-Snooping on the Distribution-Tier Switch .....	42-6
Managing RADIUS-Snooping .....	42-6
Displaying RADIUS-Snooping Statistics .....	42-6
RADIUS-Snooping Configuration Example .....	42-8
Configure the Distribution-tier Switch .....	42-9
Managing RADIUS-Snooping on the Distribution-tier Switch .....	42-9
Terms and Definitions .....	42-10

## Chapter 43: Authentication Configuration

Using Authentication in Your Network .....	43-1
Implementing User Authentication .....	43-2
Authentication Overview .....	43-2
Quarantine .....	43-2
IEEE 802.1x Using EAP .....	43-3
MAC-Based Authentication (MAC) .....	43-4
Port Web Authentication (PWA) .....	43-4
Convergence End Point (CEP) .....	43-5
Auto-Tracking .....	43-5
Multi-User And MultiAuth Authentication .....	43-6
Remote Authentication Dial-In Service (RADIUS) .....	43-9
Configuring Authentication .....	43-15

Configuring Quarantine Agent .....	43-17
Configuring IEEE 802.1x .....	43-18
Configuring MAC-based Authentication .....	43-19
Configuring Port Web Authentication (PWA) .....	43-20
Configuring Convergence End Point (CEP) .....	43-21
Configuring Auto-Tracking .....	43-22
Configuring MultiAuth Authentication .....	43-23
Configuring RADIUS .....	43-28
Authentication Configuration Example .....	43-31
Configuring the Quarantine Agent .....	43-32
Configuring the Auto-Tracking Agent .....	43-32
Setting MultiAuth Configuration On the Switch .....	43-32
Enabling RADIUS On the Switch .....	43-33
Creating RADIUS User Accounts On The Authentication Server .....	43-33
Configuring the Engineering Group 802.1x End-User Stations .....	43-33
Configuring the Engineering Group Siemens CEP Devices .....	43-34
Configuring the Printer Cluster for MAC-Based Authentication .....	43-35
Configuring the Public Area PWA Station .....	43-35
Terms and Definitions .....	43-36

## Procedures

1-1	Initial Setup .....	1-2
3-1	Executing the Configuration Restore-Point .....	3-3
3-2	Deleting the Configuration Restore-Point .....	3-4
3-3	Running a Configuration Script .....	3-8
4-1	Configuring HAU .....	4-7
5-1	Configuring VSB .....	5-11
6-1	Configuring Ports .....	6-14
6-2	Configuring Transmit Queue Monitoring .....	6-14
6-3	Configuring Link Trap and Link Flap Detection .....	6-15
7-1	Configuring OAM .....	7-9
9-1	User Management Configuration .....	9-5
9-2	Authentication Configuration .....	9-8
9-3	WebView Configuration .....	9-8
9-4	Management Authentication Notification MIB Configuration .....	9-9
9-5	License Configuration .....	9-10
9-6	Configuring SNTP .....	9-13
9-7	Telnet Configuration .....	9-17
9-8	SSH Configuration .....	9-19
9-9	Configuring DNS Resolution .....	9-21
9-10	Enabling the DHCP Server and Configuring Automatic Address Assignment .....	9-31
9-11	DHCP Client Configuration .....	9-32
9-12	Configuring DHCPv6 Information Option Pools .....	9-33
9-13	DHCPv6 Client Configuration .....	9-33
9-14	Configuring Node Alias .....	9-35
9-15	Configuring MAC Address Settings .....	9-39
11-1	Configuring an IKE Proposal .....	11-7
11-2	Configuring an IKE Policy .....	11-8
11-3	Configuring an IKE Map .....	11-8
11-4	Configuring IPsec .....	11-9
12-1	State Probe Configuration .....	12-9
12-2	Timing Probe Configuration .....	12-12
12-3	Port Group Tracked Object Configuration .....	12-13
13-1	IP SLA Configuration .....	13-6
14-1	PoE Configuration .....	14-4

15-1	Configuring LLDP (Extreme Networks 7100-Series).....	15-9
17-1	New SNMPv1/v2c Configuration .....	17-9
17-2	SNMPv3 Configuration .....	17-11
17-3	Configuring an EngineID .....	17-14
17-4	Configuring Secure Community Names .....	17-18
19-1	Configuring Shortest Path Bridging .....	19-6
20-1	Static VLAN Configuration .....	20-11
20-2	Secure Management VLAN Configuration .....	20-14
20-3	GVRP Configuration .....	20-14
20-4	MVRP Configuration .....	20-15
20-5	Configuring Protocol-Based VLAN Classification .....	20-15
20-6	IGMP Snooping for a VLAN Configuration .....	20-17
20-7	Configuring a Provider Bridge .....	20-22
21-1	Configuring Link Aggregation .....	21-9
22-1	Configuring Policy Roles .....	22-11
22-2	Configuring Classification Rules .....	22-12
23-1	Basic IGMP Configuration .....	23-21
23-2	Basic DVMRP Configuration .....	23-23
23-3	Basic PIM Sparse Mode Configuration .....	23-28
24-1	Basic MLD Configuration .....	24-6
25-1	Configuring a Server and Console Logging .....	25-12
25-2	Adjusting Settings for an Application .....	25-12
26-1	Configuring SMON .....	26-9
26-2	Configuring Remote Network Monitoring .....	26-10
27-1	CFM Maintenance Domain (MD) Configuration .....	27-26
27-2	CFM Maintenance Association (MA) Configuration .....	27-27
27-3	CFM Maintenance Association Component (MA-Comp) Configuration .....	27-27
27-4	CFM Maintenance Association End-Point (MEP) Configuration .....	27-28
28-1	VRF Configuration .....	28-4
29-1	Configuring the Routing Interface .....	29-12
29-2	Configuring Non-forward IP Static Routes .....	29-13
29-3	Configuring an IPv6 Static Neighbor Discovery Cache Entry .....	29-15
29-4	Configuring the ARP Table .....	29-17
30-1	Configuring RIP .....	30-4
31-1	Configuring RIPng .....	31-4
32-1	Configuring Basic OSPF Parameters .....	32-23
32-2	Configuring OSPF General Optional Parameters .....	32-24
32-3	Configuring OSPF Optional Interface Parameters .....	32-26
33-1	Configuring Basic OSPFv3 Parameters .....	33-27
34-1	Configuring Global IS-IS .....	34-15
34-2	Configuring IS-IS IPv6 Unicast Address Family .....	34-18
35-1	Configuring Basic BGP .....	35-25
35-2	Configuring BGP Route Injection .....	35-26
35-3	EBGP Basic Peering Configuration .....	35-28
35-4	IBGP Basic Peering Configuration .....	35-31
35-5	Multihop BGP Basic Peering Configuration .....	35-33
35-6	Configuring Source IP Address to the Peer Update .....	35-36
35-7	Configuring BGP Confederation .....	35-39
35-8	Configuring BGP Route Reflection .....	35-42
35-9	Configuring BGP Conditional Route Advertisement .....	35-46
35-10	Configuring BGP Flap Dampening .....	35-49
35-11	Configuring Graceful Restart .....	35-50
36-1	Configuring VRRP .....	36-8
37-1	MAC Locking Configuration .....	37-8
37-2	SSH Configuration .....	37-10
37-3	TACACS+ Configuration .....	37-11



37-4	Host DoS Configuration .....	37-12
38-1	Creating and Managing IPv4 and IPv6 ACLs .....	38-9
38-2	Entering and Managing Standard IPv4 ACL Rules.....	38-9
38-3	Entering and Managing Standard IPv6 ACL Rules.....	38-9
38-4	Entering and Managing Extended IPv4 ACL Rules.....	38-10
38-5	Entering and Managing Extended IPv6 ACL Rules.....	38-12
38-6	Managing IPv4 and IPv6 ACL Rules .....	38-13
38-7	Applying and Displaying ACLs .....	38-14
39-1	Configuring a Policy Based Route-Map.....	39-8
39-2	Configuring a Redistribution Route-Map .....	39-9
39-3	Configuring a Filter Route-Map .....	39-10
39-4	Configuring a BGP Route-Map.....	39-11
40-1	Class of Service CLI Configuration Command Summary.....	40-17
41-1	Configuring Anti-Spoofing Features.....	41-7
42-1	RADIUS-Snooping Configuration .....	42-6
43-1	Quarantine Agent Configuration .....	43-17
43-2	IEEE 802.1x Configuration .....	43-18
43-3	MAC-Based Authentication Configuration .....	43-19
43-4	Port Web Authentication (PWA) Configuration .....	43-20
43-5	CEP Detection Group Configuration.....	43-21
43-6	CEP Configuration.....	43-21
43-7	Configuring MultiAuth Idle and Session Timeouts for CEP .....	43-22
43-8	Auto-tracking Agent Configuration.....	43-23
43-9	MultiAuth Authentication Configuration .....	43-24
43-10	MultiAuth Authentication Precedence Configuration .....	43-24
43-11	MultiAuth Authentication Port and Maximum User Properties Configuration .....	43-25
43-12	MultiAuth Authentication Timers Configuration.....	43-26
43-13	MultiAuth Authentication Traps Configuration .....	43-26
43-14	VLAN Authorization Configuration.....	43-28
43-15	Policy Profile Assignment and Invalid Action Configuration .....	43-28
43-16	Authentication Server Configuration.....	43-29
43-17	RADIUS Accounting Configuration.....	43-30

## Figures

4-1	VSF System High Availability Firmware Upgrade .....	4-6
5-1	VSF Stacking Data center Configuration Overview.....	5-2
5-2	VSF LFR Stack Examples.....	5-8
7-1	Frame Link Monitor Option .....	7-4
7-2	Frame-Seconds Link Monitor Option.....	7-5
7-3	Frame-Period Link Monitor Option .....	7-5
7-4	Symbol-Period Link Monitor Option.....	7-7
8-1	Using Port Mirroring to Monitor a Departmental Switch .....	8-2
8-2	Using Port Mirroring to Monitor Incoming Traffic to a Backbone Switch .....	8-2
15-1	Communication between LLDP-enabled Devices .....	15-3
15-2	LLDP-MED .....	15-5
15-3	Frame Format.....	15-6
16-1	Enhanced Transmission Selection (ETS) Queuing .....	16-3
16-2	Congestion Notification Overview.....	16-6
16-3	Congestion Notification Domain Defense Mode Overview.....	16-10
18-1	Redundant Link Causes a Loop in a Non-STP Network .....	18-2
18-2	Loop Avoided When STP Blocks a Duplicate Path .....	18-2
18-3	Multiple Spanning Tree Overview.....	18-5
18-4	Root Port Selection Based Upon Lowest Cost or Bridge ID.....	18-11
18-5	Root Port Selection Base Upon Lowest Port ID .....	18-12
18-6	Spanning Tree Port Role Overview .....	18-13

18-7	Example of an MST Region.....	18-16
18-8	MSTI 1 in a Region.....	18-19
18-9	MSTI 2 in the Same Region .....	18-19
18-10	Example of Multiple Regions and MSTIs.....	18-20
18-11	Traffic Segregation in a Single STP Network Configuration.....	18-28
18-12	Traffic Segregation in an MSTP Network Configuration.....	18-29
18-13	Maximum Bandwidth Utilization in a Single STP Network Configuration .....	18-30
18-14	Maximum Bandwidth Utilization in an MSTP Network Configuration .....	18-31
18-15	Basic Loop Protect Scenario .....	18-36
18-16	Spanning Tree Without Loop Protect .....	18-36
18-17	Spanning Tree with Loop Protect .....	18-36
19-1	SPB Overview .....	19-2
19-2	SPBV Using Equal Cost Trees .....	19-3
20-1	VLAN Business Scenario .....	20-2
20-2	Inside the Switch .....	20-6
20-3	Example of VLAN Propagation Using GVRP or MVRP.....	20-9
20-4	Provider Bridges in Provider Network.....	20-21
20-5	Provider Bridge Configuration Example .....	20-23
21-1	LAG Formation .....	21-4
21-2	LAGs Moved to Attached State .....	21-6
21-3	Example 1 Multiple Device Configuration.....	21-12
21-4	Example 2 Configuration .....	21-17
22-1	College-Based Policy Configuration .....	22-15
23-1	IGMP Querier Determining Group Membership .....	23-3
23-2	Sending a Multicast Stream with No Directly Attached Hosts .....	23-4
23-3	DVMRP Pruning and Grafting .....	23-10
23-4	PIM Traffic Flow.....	23-12
23-5	Anycast-RP Configuration .....	23-17
23-6	PIM-DM Traffic Flow .....	23-18
23-7	DVMRP Configuration on Two Routers.....	23-24
23-8	PIM-SM Configuration with Bootstrap Router and Candidate RPs .....	23-30
23-9	PIM-SSM Configuration.....	23-33
24-1	MLD Querier Determining Group Membership.....	24-3
24-2	Sending a Multicast Stream with No Directly Attached Hosts .....	24-4
25-1	Basic Syslog Scenario.....	25-5
27-1	Maintenance Domain Overview.....	27-3
27-2	Maintenance Association Overview.....	27-5
27-3	Maintenance End-Point Overview .....	27-6
27-4	Maintenance Intermediate-Point Overview.....	27-7
27-5	VLAN Table Configuration Overview.....	27-10
27-6	Single MD Example Configuration Overview.....	27-32
27-7	Multiple MD Example Configuration Overview .....	27-40
28-1	VRF Overview .....	28-3
29-1	Secondary VLAN Configuration.....	29-8
32-1	Basic OSPF Topology .....	32-5
32-2	OSPF Router ID Topology.....	32-6
32-3	OSPF Designated Router Topology .....	32-8
32-4	OSPF Summarization Topology .....	32-10
32-5	OSPF Stub Area Topology .....	32-12
32-6	OSPF NSSA Topology .....	32-14
32-7	Virtual Link Topology .....	32-15
32-8	Physical and Logical Single Router HA Failover Configuration.....	32-18
33-1	Basic OSPF Topology .....	33-7
33-2	OSPF Designated Router Topology .....	33-10
33-3	OSPF Summarization Topology .....	33-12
33-4	OSPF Stub Area Topology .....	33-14

33-5	OSPF NSSA Topology .....	33-15
33-6	Virtual-Link Topology .....	33-18
33-7	Physical and Logical Single Router HA Failover Configuration .....	33-22
34-1	IS-IS Network Overview .....	34-2
34-2	Network Layer Addresses (NSAP) .....	34-4
34-3	IS-IS NET Configuration .....	34-7
34-4	IS-IS Route Summarization .....	34-10
35-1	BGP Topology .....	35-3
35-2	Route Flap Dampening Timing .....	35-18
35-3	Basic EBGP Peering Topology .....	35-27
35-4	Basic IBGP Peering Topology .....	35-29
35-5	EBGP Multihop Peering Topology .....	35-31
35-6	Source IP Address to a Remote Peer .....	35-36
35-7	BGP Confederation Example Topology .....	35-37
35-8	BGP Route Reflection Example Topology .....	35-40
35-9	BGP Conditional Advertisement Example Topology .....	35-43
35-10	Route Flap Dampening Example Configuration .....	35-46
36-1	A Basic VRRP Topology .....	36-2
36-2	Critical-IP Address Configuration .....	36-4
36-3	Standard VRRP Forwarding .....	36-6
36-4	Fabric Route VRRP Forwarding .....	36-7
36-5	Basic Configuration Example .....	36-10
36-6	Multi-Backup VRRP Configuration Example .....	36-11
37-1	Blocking Unauthorized Access with MAC Locking .....	37-4
40-1	Assigning and Marking Traffic with a Priority .....	40-3
40-2	Strict Priority Queuing Packet Behavior .....	40-4
40-3	Weighted Fair Queuing Packet Behavior .....	40-5
40-4	Hybrid Queuing Packet Behavior .....	40-6
40-5	Rate Limiting Clipping Behavior .....	40-7
40-6	Rate Shaping Smoothing Behavior .....	40-7
42-1	RADIUS-Snooping Overview .....	42-4
42-2	RADIUS-Snooping Configuration Example Overview .....	42-8
43-1	Applying Policy to Multiple Users on a Single Port .....	43-7
43-2	Authenticating Multiple Users With Different Methods on a Single Port .....	43-8
43-3	Selecting Authentication Method When Multiple Methods are Validated .....	43-9
43-4	Authentication Configuration Example Overview .....	43-31

## Tables

1-1	Advanced Configuration .....	1-2
2-1	CLI Properties Configuration Commands .....	2-4
2-2	CLI Properties Show Commands .....	2-5
3-1	Configuration and Image File Management and Display Commands .....	3-9
4-1	Default HAU Parameters .....	4-6
4-2	HAU Configuration Terms and Definitions .....	4-8
5-1	Default VSB Parameters .....	5-11
5-2	VSB Configuration Terms and Definitions .....	5-12
6-1	Default Port Parameters .....	6-12
6-2	Managing Port Configuration .....	6-16
6-3	Displaying Port Configuration Information and Statistics .....	6-16
6-4	Port Configuration Terms and Definitions .....	6-17
7-1	Frame-Period Window Values .....	7-5
7-2	Symbol-Period Window Values .....	7-6
7-3	Default Ethernet OAM Configuration Settings .....	7-8
7-4	OAM Configuration Terms and Definitions .....	7-11
9-1	Default System Parameters .....	9-1

9-2	System Properties Configuration .....	9-2
9-3	System Properties Management and Display Commands .....	9-3
9-4	User Account Management and Display Commands .....	9-6
9-5	Default SNMP Parameters .....	9-13
9-6	Managing and Displaying SNMP .....	9-14
9-7	Default DNS Parameters .....	9-21
9-8	Managing DNS Resolution .....	9-22
9-9	IPv4 DHCP Server Codes .....	9-24
9-10	DHCPv6 Server Supported Options .....	9-29
9-11	Default DHCP Parameters .....	9-31
9-12	Configuring Static IP Address Assignment .....	9-32
9-13	Managing and Displaying DHCP .....	9-34
9-14	Managing Node Alias .....	9-35
9-15	System Configuration Terms and Definitions .....	9-39
10-1	Security Profile Mode Command Parameter Default Setting Changes .....	10-3
10-2	Security Profile Mode Command Parameter Range Changes .....	10-4
10-3	Security Profile mode Command Access Changes .....	10-4
10-4	Read-Write Functionality Not Accessible in C2 Security Profile Mode .....	10-4
10-5	Read-Only Functionality Not Accessible in C2 Security Profile Mode .....	10-6
10-6	Configuring Security Mode on the Device .....	10-6
10-7	Security Mode Show Commands .....	10-7
10-8	Security Mode Configuration Terms and Definitions .....	10-7
11-1	IKE Proposal Parameters .....	11-4
11-2	IKE Policy Parameters .....	11-5
11-3	IPsec Show Commands .....	11-10
11-4	IPsec Configuration Terms and Definitions .....	11-12
12-1	Default Tracked Object Manager Parameters .....	12-8
12-2	Configuring Timing Probe Parameters .....	12-10
12-3	Configuring Tracked Object Parameters .....	12-13
12-4	Tracked Object Manager Terms and Definitions .....	12-15
13-1	Default IP SLA Values .....	13-5
13-2	IP SLA Show Commands .....	13-7
14-1	PoE Powered Device Classes .....	14-2
14-2	Default PoE Parameter Values .....	14-3
14-3	PoE Show Commands .....	14-7
15-1	LLDP Configuration Commands .....	15-7
15-2	LLDP Show Commands .....	15-10
15-3	Neighbor Warning Detection Commands .....	15-12
15-4	Enterasys Discovery Protocol Configuration Commands .....	15-12
15-5	Enterasys Discovery Protocol Show Commands .....	15-13
15-6	Cisco Discovery Protocol Configuration Commands .....	15-13
15-7	Cisco Discovery Protocol Show Commands .....	15-14
16-1	Choice, Defense Mode, and Alternate Priority Cross-Reference .....	16-12
16-2	Default Data Center Bridging Configuration Settings .....	16-14
16-3	Data Center Bridging Configuration .....	16-15
16-4	Congestion Notification Global Configuration .....	16-15
16-5	Congestion Notification Global Priority Configuration .....	16-16
16-6	Congestion Notification Port Priority Configuration .....	16-16
16-7	Data Center Bridging Display Commands .....	16-17
16-8	Data Center Bridging (DCB) Configuration Terms and Definitions .....	16-17
17-1	SNMP Message Functions .....	17-3
17-2	SNMP Terms and Definitions .....	17-5
17-3	SNMP Security Models and Levels .....	17-7
17-4	Default Extreme Networks SNMP Configuration .....	17-9
18-1	Spanning Tree Port Roles .....	18-14
18-2	Spanning Tree Port States .....	18-14

18-3	MSTI Characteristics for Figure 18-10.....	18-20
18-4	Spanning Tree Port Default Settings .....	18-22
18-5	BPDU Interval Defaults.....	18-24
18-6	Commands for Monitoring MSTP .....	18-32
18-7	Commands for Monitoring SpanGuard.....	18-34
18-8	Commands for Monitoring Loop Protect.....	18-38
18-9	Spanning Tree Terms and Definitions .....	18-39
19-1	Default Shortest Path Bridging Parameters.....	19-6
19-2	Shortest Path Bridging Configuration Terms and Definitions .....	19-7
20-1	Default VLAN Parameters .....	20-10
20-2	Displaying VLAN Information.....	20-18
20-3	VLAN Terms and Definitions .....	20-18
20-4	Provider Bridge VLAN Tags .....	20-20
20-5	Provider Bridge-related Port Types .....	20-21
21-1	LAG2 Port Priority Assignments .....	21-5
21-2	LAG Port Parameters .....	21-7
21-3	Extreme Networks Platform LAG Support.....	21-8
21-4	Default Link Aggregation Parameters.....	21-9
21-5	Managing Link Aggregation.....	21-10
21-6	Displaying Link Aggregation Information and Statistics.....	21-11
21-7	LAG and Physical Port Admin Key Assignments .....	21-13
21-8	Link Aggregation Configuration Terms and Definitions .....	21-19
22-1	Administrative Policy and Policy Rule Traffic Classifications .....	22-7
22-2	Non-Edge Protocols .....	22-8
22-3	Policy Traffic Classification Sets.....	22-10
22-4	Traffic Classification Based Policy Capabilities .....	22-10
22-5	Displaying Policy Configuration and Statistics.....	22-13
22-6	Policy Configuration Terms and Definitions.....	22-25
23-1	PIM Terms and Definitions .....	23-19
23-2	IGMP Configuration Commands.....	23-20
23-3	Layer 2 IGMP Show Commands .....	23-22
23-4	Layer 3 IGMP Show Commands .....	23-22
23-5	DVMRP Configuration Commands.....	23-23
23-6	DVMRP Show Commands .....	23-24
23-7	IPv4 PIM Sparse Mode Commands .....	23-25
23-8	IPv6 PIM Sparse Mode Commands .....	23-26
23-9	PIM Dense Mode Commands .....	23-27
23-10	PIM IPv4 and IPv6 Display Commands.....	23-28
24-1	MLD Configuration Commands .....	24-5
24-2	MLD Show Commands.....	24-6
25-1	Syslog Terms and Definitions .....	25-3
25-2	Syslog Message Components.....	25-6
25-3	Syslog Command Precedence .....	25-7
25-4	Syslog Server Default Settings.....	25-8
26-1	RMON Monitoring Group Functions and Commands.....	26-5
26-2	Default Network Monitoring Parameters.....	26-7
26-3	Network Diagnostics Commands .....	26-8
26-4	Managing Network Monitoring.....	26-11
26-5	Displaying Network Monitoring Information and Statistics.....	26-12
27-1	CFM Configuration Modes.....	27-8
27-2	MD Configuration Modes.....	27-12
27-3	MEP Defect Definitions.....	27-21
27-4	Default Connectivity Fault Management Configuration Settings .....	27-24
27-5	CFM Global Configuration .....	27-25
27-6	CFM Global and Monitored VLAN Service Default Configuration .....	27-26
27-7	CFM Management Commands .....	27-28

27-8	CFM Show Commands .....	27-29
27-9	Connectivity Fault Management (CFM) Terms and Definitions .....	27-44
28-1	Default VRF Parameters .....	28-4
28-2	VRF Configuration Terms and Definitions .....	28-4
29-1	Entering Router Configuration Mode .....	29-2
29-2	Default IP Routing Parameters .....	29-18
29-3	Managing the Router .....	29-19
29-4	Displaying IP Routing Information and Statistics .....	29-19
29-5	IP Routing Terms and Definitions .....	29-20
30-1	Default RIP Parameters .....	30-4
30-2	RIP Configuration Terms and Definitions .....	30-5
31-1	Default RIPng Parameters .....	31-3
31-2	RIPng Configuration Terms and Definitions .....	31-4
32-1	Default OSPF Parameters .....	32-22
32-2	Displaying OSPF Configuration and Statistics .....	32-27
33-1	OSPFv3 and OSPFv2 LSA Cross-Reference .....	33-3
33-2	Default OSPF Parameters .....	33-25
33-3	Configuring OSPFv3 General Optional Parameters .....	33-27
33-4	Configuring OSPF Optional Interface Parameters .....	33-29
33-5	Displaying OSPFv3 Configuration and Statistics .....	33-30
34-1	IS-IS Timers .....	34-11
34-2	IS-IS Parameters .....	34-14
34-3	Configuring IS-IS on the Interface .....	34-17
34-4	Displaying IS-IS Information .....	34-18
34-5	Feature Configuration Terms and Definitions .....	34-18
35-1	AS-Path Regular Expressions .....	35-7
35-2	Default BGP Parameters .....	35-23
35-3	BGP Neighbor Configuration .....	35-34
35-4	Configuring BGP Outbound Route Filtering .....	35-43
35-5	Configuring BGP Soft Reset .....	35-46
35-6	Monitoring and Clearing BGP Configuration .....	35-51
35-7	BGP Terms and Definitions .....	35-51
36-1	Default VRRP Parameters .....	36-8
36-2	Displaying VRRP Information and Statistics .....	36-9
36-3	VRRP Configuration Terms and Definitions .....	36-13
37-1	Host DoS Mitigation Types .....	37-7
37-2	Default Security Parameters .....	37-8
37-3	Managing MAC Locking .....	37-9
37-4	Managing TACACS+ .....	37-11
37-5	Displaying Host DoS .....	37-13
38-1	ACL Configuration Terms and Definitions .....	38-14
39-1	Default Route-Map Manager Parameters .....	39-8
39-2	Displaying Route-Map Manager Information and Statistics .....	39-13
39-3	Route-Map Manager Terms and Definitions .....	39-15
40-1	Quality of Service Configuration Terms and Definitions .....	40-18
41-1	Managing Anti-Spoofing Features .....	41-8
41-2	Displaying Anti-Spoofing Information .....	41-8
42-1	Default Authentication Parameters .....	42-5
42-2	Managing RADIUS-Snooping .....	42-6
42-3	Displaying RADIUS-Snooping Statistics .....	42-6
42-4	RADIUS-Snooping Configuration Terms and Definitions .....	42-10
43-1	Default Authentication Parameters .....	43-15
43-2	MultiAuth Authentication Settings and Statistics Display .....	43-27
43-3	Quality of Service Configuration Terms and Definitions .....	43-36



---

# About This Guide

This manual explains how to configure Extreme Networks 7100-Series switch/router devices.

## How to Use This Guide

Read through this guide completely to familiarize yourself with its contents and to gain an understanding of the features and capabilities of the 7100-Series modules. A general working knowledge of data communications networks is helpful when setting up these modules.

## Related Documents

The manuals listed below can be obtained from the World Wide Web in Adobe Acrobat Portable Document Format (PDF) at the following site:

<http://support.extremenetworks.com/>

- *Extreme Networks 7100-Series CLI Reference* provides information on how to use the Command Line Interface for the 7100-Series switch/routers.

## Conventions Used in This Guide

The following conventions are used in the text of this document:

Convention	Description
<b>Bold font</b>	Indicates mandatory keywords, parameters or keyboard keys.
<i>italic font</i>	Indicates complete document titles.
Courier font	Used for examples of information displayed on the screen.
<i>Courier font in italics</i>	Indicates a user-supplied value, either required or optional.
[ ]	Square brackets indicate an optional value.
{ }	Braces indicate required values. One or more values may be required.
	A vertical bar indicates a choice in values.
[x   y   z]	Square brackets with a vertical bar indicates a choice of a value.
{x   y   z}	Braces with a vertical bar indicate a choice of a required value.
[x {y   z}]	A combination of square brackets with braces and vertical bars indicates a required choice of an optional value.

The following icons are used in this guide:



**Note:** Calls the reader's attention to any item of information that may be of special importance.





**Router:** Calls the reader's attention to router-specific configuration information.



**Caution:** Contains information essential to avoid damage to the equipment.

**Precaución:** Contiene información esencial para prevenir dañar el equipo.

**Achtung:** Verweist auf wichtige Informationen zum Schutz gegen Beschädigungen.

## Commonly Used Acronyms

The following acronyms are used extensively throughout this guide:

- LED – Light Emitting Diode
- USB – Universal Serial Bus

## Getting Help

For additional support related to 7100-Series switch/router or to this document, contact Enterasys Networks using one of the following methods:

World Wide Web	<a href="http://support.extremenetworks.com/">http://support.extremenetworks.com/</a>
	1-800-872-8440 (toll-free in U.S. and Canada) or 1-603-952-5000
	For the Extreme Networks Support toll-free number in your country:
Phone	<a href="http://www.extremenetworks.com/support/contact/">www.extremenetworks.com/support/contact/</a> <a href="mailto:support@extremenetworks.com">support@extremenetworks.com</a>
Internet mail	To expedite your message, please type <b>[7100-Series]</b> in the subject line.

**Before contacting Enterasys Networks for technical support, have the following data ready:**

- Your Enterasys Networks service contract number
- A description of the failure
- A description of any action(s) already taken to resolve the problem (for example, changing mode switches or rebooting the unit)
- The serial and revision numbers of all involved Enterasys Networks products in the network
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load and frame size at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any previous Return Material Authorization (RMA) numbers

## Getting Started

This chapter provides the procedures to start the 7100-Series device once the hardware is installed. Initially, the system can only be configured using the Command Line Interface (CLI) from a device connected directly to the console port on the chassis.

This chapter also provides an overview of configuring the 7100-Series as a switch and router to fit into your network.

For information about...	Refer to page...
<a href="#">Device Management Methods</a>	1-1
<a href="#">Initial Configuration</a>	1-1
<a href="#">Advanced Configuration Overview</a>	1-2



**Notes:** See the default parameters table located in the relevant chapter for factory default values.

## Device Management Methods

The 7100-Series device can be managed using the following methods:

- Locally using a VT type terminal connected to the console port.
- Remotely using a VT type terminal connected through a modem.
- Remotely using an SNMP management station.
- In-band through a Telnet connection.
- In-band using Extreme Networks' NetSight management application.
- Remotely using WebView™, Extreme Networks' embedded web server application.

The *Hardware Installation Guide* for your 7100-Series device provides setup instructions for connecting a terminal or modem to the device.

## Initial Configuration

To initially configure the 7100-Series device, you must have connected a terminal to the local console port as described in the *Hardware Installation Guide* for your 7100-Series device. [Procedure 1-1](#) contains the steps to assign an IP address and configure basic system parameters. For information on the command syntax and parameters, refer to the online help or the *Extreme Networks 7100-Series CLI Reference*.



**Note:** When configuring any string or name parameter input for any command, do not use any letters with diacritical marks (an ancillary glyph added to a letter). Diacritical marked letters are not supported by SNMP.

### Procedure 1-1 Initial Setup

Step	Task	Command
1.	Log in as an administrator.	<ul style="list-style-type: none"> <li>At the login prompt, enter <b>admin</b>.</li> <li>Press <b>Enter</b> for the password (no password string by default).</li> </ul>
2.	For security, change the password.	<b>set password</b>
3.	Optionally, check the version of the firmware image then check the Extreme Networks web site to verify that you have the latest version.	<b>show version</b>
4.	Optionally, define a name for the system, the location of the system, and contact information for system issues.	<b>set system name</b> <i>[string]</i> <b>set system location</b> <i>[string]</i> <b>set system contact</b> <i>[string]</i>
5.	Optionally, define a message that displays whenever a user logs in.	<b>set banner</b> { <b>motd</b>   <b>login</b> } <i>message</i>
6.	Optionally, change the default prompt.	<b>set prompt</b> " <i>prompt_string</i> "
7.	Display the system's setting for the date and time. If necessary, change the setting.  <b>NOTE:</b> Instead of manually setting the time, you can configure the system as an SNTP client, as described in " <a href="#">SNTP Overview</a> " on page 9-11.	<b>show time</b> <b>set time</b> <i>[mm/dd/yyyy] [hh:mm:ss]</i>
8.	Assign a management IP address.	<b>set ip interface</b> <b>set ip address</b>
9.	If desired, configure additional user accounts and passwords. Up to 32 user accounts may be registered with the local database.	<b>set system login</b> <i>username</i>

## Advanced Configuration Overview

The 7100-Series device can be configured to provide various system services, Layer 2 switching, Layer 3 routing, and security. [Table 1-1](#) provides an overview of configuring the 7100-Series device for each area.



**Note:** Though it is possible to configure policy by using the CLI, Extreme Networks recommends that you use NetSight Policy Manager instead.

**Table 1-1 Advanced Configuration**

Task	Refer to page...
<b>System Services</b>	
Configure the Simple Network Time Protocol (SNTP) client.	<a href="#">9-11</a>
Configure the Domain Name Server (DNS) client.	<a href="#">9-21</a>

**Table 1-1 Advanced Configuration (continued)**

Task	Refer to page...
Configure the Telnet client and server. (Telnet client is enabled by default.) <b>Note:</b> For security, you may wish to disable Telnet and only use SSH.	<a href="#">9-16</a>
Configure the Secure Shell V2 (SSHv2) client and server.	<a href="#">9-17</a>
Configure the Dynamic Host Configuration Protocol (DHCP) client and server.	<a href="#">9-23</a>
Configure the port parameters, such as speed and duplex mode.	<a href="#">6-1</a>
Enable SNMP and create a community string. By default, the SNMP master agent is disabled and no defined public community string is configured.	<a href="#">17-18</a>
Configure RMON to provide comprehensive network fault diagnosis, planning, and performance tuning information, and allow for interoperability between SNMP management stations and monitoring agents.	<a href="#">26-4</a>
Change the interactive login authentication method, from local to remote (RADIUS authentication).	<a href="#">43-28</a>
If RADIUS authentication is configured, configure the remote RADIUS servers to be used by the RADIUS client on the 7100-Series	<a href="#">43-28</a>
<b>Layer 2 Switching</b>	
Enable desired ports for switching.	<a href="#">6-4</a>
Set port configurations and port-based Virtual Local Area Networks (VLANs). VLANs can be created statically or dynamically.	<a href="#">24-3</a>
Configure Spanning Trees using STP, RSTP, or MSTP.	<a href="#">18-21</a>
Configure LLDP or CDP.	<a href="#">15-1</a>
<b>Layer 3 Routing</b>	
Configure the router id. Refer to the <b>router id</b> command in the <i>Extreme Networks 7100-Series CLI Reference</i> .	
Configure interfaces for IP routing.	<a href="#">29-3</a>
Configure the ARP table.	<a href="#">29-15</a>
Configure routes.	<a href="#">29-12</a>
Configure multicast protocols IGMP parameters.	<a href="#">23-20</a>
<b>Security and General Management</b>	
Configure RADIUS servers.	<a href="#">43-28</a>
Manage user accounts and passwords.	<a href="#">1-1</a>
Configure system logging.	<a href="#">25-6</a>
Configure the 7100-Series using text files.	<a href="#">3-1</a>
Upgrade system firmware.	<a href="#">3-1</a>
Configure QoS features.	<a href="#">40-8</a>
Configure policy.	<a href="#">22-10</a>



## Using the CLI

This chapter provides information about CLI conventions for 7100-Series devices and CLI properties that you can configure.

For information about...	Refer to page...
<a href="#">CLI Conventions</a>	2-1
<a href="#">Configuring CLI Properties</a>	2-4

### CLI Conventions

For information about...	Refer to page...
<a href="#">Getting Help with CLI Syntax</a>	2-1
<a href="#">Using Context-Sensitive Help</a>	2-1
<a href="#">Performing Keyword Lookups</a>	2-2
<a href="#">Displaying Scrolling Screens</a>	2-3
<a href="#">Abbreviating and Completing Commands</a>	2-3
<a href="#">Using the Spacebar Auto Complete Function</a>	2-4

### Getting Help with CLI Syntax

The 7100-Series device allows you to display usage and syntax information for individual commands by typing **help** or **?** after the command.

### Using Context-Sensitive Help

Entering **help** after a specific command will display usage and syntax information for that command. This example shows how to display context-sensitive help for the **set length** command:

```
7100 Chassis(rw)->set length help
Command: set length Number of lines
Usage: set length <screenlength>
       screenlength      Length of the screen (5..512, 0 to disable 'more')
```

## Performing Keyword Lookups

Entering a space and a question mark (?) after a keyword will display all commands beginning with the keyword. The following example shows how to perform a keyword lookup for the **show snmp** command. In this case, 13 additional keywords are used by the **show snmp** command. Entering a space and a question mark (?) after any of these parameters (such as **show snmp user**) will display additional parameters nested within the syntax.

```
7100 Chassis(rw)->show snmp ?
access          SNMP VACM access configuration
community      SNMP v1/v2c community name configuration
context        SNMP VACM context list
counters       SNMP counters
engineid       SNMP engine properties
group          SNMP VACM security to group configuration
notify         SNMP notify configuration
notifyfilter   SNMP notify filter configuration
notifyprofile  SNMP notify profile configuration
targetaddr     SNMP target address configuration
targetparams   SNMP target parameters configuration
user           SNMP USM user configuration
view           SNMP VACM view tree configuration

7100 Chassis(rw)->show snmp
7100 Chassis(rw)->show snmp user ?
list           List usernames
<user>        User name
remote        Show users with remote SNMP engine ID
volatile      Show temporary entries
nonvolatile   Show permanent entries
read-only     Show r/o entries
<cr>

7100 Chassis(rw)->show snmp user
```

Entering a question mark (?) without a space after a partial keyword will display a list of commands that begin with the partial keyword. The following example shows how to use this function for all commands beginning with **co**:

```
7100 Chassis(rw)->co?
configure      Execute a configuration file
copy          Upload or download an image or configuration file

7100 Chassis(rw)->co
```



**Note:** At the end of the lookup display, the system will repeat the command you entered without the ?.

## Displaying Scrolling Screens

If the CLI screen length has been set using the **set length** command as described in [Table 2-1](#) on page 2-4, CLI output requiring more than one screen will display `--More--` to indicate continuing screens. To display additional screen output:

- Press any key other than ENTER to advance the output one screen at a time.
- Press ENTER to advance the output one line at a time.

The following example shows how the **show mac** command indicates that output continues on more than one screen.

```
7100 Chassis(rw)->show mac
```

MAC Address	FID	Port	Type
00-00-1d-67-68-69	1	host.0.1	learned
00-00-02-00-00-00	1	ge.1.2	learned
00-00-02-00-00-01	1	ge.1.3	learned
00-00-02-00-00-02	1	ge.1.4	learned
00-00-02-00-00-03	1	ge.1.5	learned
00-00-02-00-00-04	1	ge.1.6	learned
00-00-02-00-00-05	1	ge.1.7	learned
00-00-02-00-00-06	1	ge.1.8	learned
00-00-02-00-00-07	1	ge.1.9	learned
00-00-02-00-00-08	1	ge.1.10	learned

--More--

## Abbreviating and Completing Commands

The 7100-Series device allows you to abbreviate CLI commands and keywords down to the number of characters that will allow for a unique abbreviation. The following example shows how to abbreviate the **show netstat** command to **show net**.

```
7100 Chassis(rw)->show net
```

Active Internet connections (including servers)

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	State
TCP	0	0	10.21.73.13.23	134.141.190.94.51246	ESTABLISHED
TCP	0	275	10.21.73.13.23	134.141.192.119.4724	ESTABLISHED
TCP	0	0	*.80	*.*	LISTEN
TCP	0	0	*.23	*.*	LISTEN
UDP	0	0	10.21.73.13.1030	134.141.89.113.514	
UDP	0	0	*.161	*.*	
UDP	0	0	*.1025	*.*	
UDP	0	0	*.123	*.*	



## Using the Spacebar Auto Complete Function

When the spacebar auto complete function is enabled, pressing the spacebar after a CLI command fragment will allow you to determine if the fragment is unique. If it is, the CLI will complete the fragment on the current display line.

By default, this function is disabled. For more information on enabling it using the **set cli completion** command, refer to [Table 2-1](#) on page 2-4. The following example shows how, when the function is enabled, entering **conf** and pressing the spacebar would be completed as **configure**:

```
7100 Chassis(rw) ->conf<SPACEBAR>
7100 Chassis(rw) ->configure
```

## Configuring CLI Properties

CLI properties are options that you can configure and customize in the CLI, such as the command prompt, command completion, banner messages, and session idle timeout.

[Table 2-1](#) lists CLI properties configuration commands.

**Table 2-1 CLI Properties Configuration Commands**

Task	Command
Modify the command prompt	<b>set prompt</b> <i>prompt-string</i>
Enable or disable the CLI command completion function. When enabled, this allows you to complete a unique CLI command fragment using the keyboard spacebar.	<b>set cli completion</b> {enable   disable} [default]
Set the banner message for pre and post session login.	<b>set banner</b> {login <i>message</i>   motd <i>message</i> }
Clear the banner message displayed at pre and post session login to a blank string.	<b>clear banner</b> {login   motd}
Set the number of columns for the terminal connected to the device's console port.	<b>set width</b> <i>screenwidth</i> [default]
Set the number of lines the CLI will display.	<b>set length</b> <i>screenlength</i> [default]
Set the time (in minutes) an idle console or Telnet CLI session will remain connected before timing out.	<b>set logout</b> <i>timeout</i> [default]
Set the current and default line editing mode or the way the Delete character is treated by the line editor. You can also set the persistence of your line editing selections.	<b>set line-editor</b> {emacs   vi   default   delete {backspace   delete}} [default]

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

## Example CLI Properties Configuration

In this example, the prompt is changed and a login banner is added.

```
7100 Chassis(rw) ->set prompt "Switch 1"
Switch 1(rw) ->
Switch 1(rw) ->set banner login There is nothing more important than our customers
```

## CLI Properties Display Commands

Table 2-2 lists CLI properties show commands.

**Table 2-2 CLI Properties Show Commands**

Task	Command
Display the current and default line-editor mode and Delete character mode.	<b>show line-editor</b>
Display the banner message that will display at pre and post session login.	<b>show banner</b>
Display the number of columns for the terminal connected to the device's console port.	<b>show width</b>
Display the current screen length.	<b>show length</b>
Display the time (in seconds) an idle console or Telnet CLI session will remain connected before timing out.	<b>show logout</b>

Refer to the *Extreme Networks 7100-Series CLI Reference* for a description of the output of each command.



## Image Configuration and File Management

This chapter provides information about configuration and image file management on the 7100-Series devices.

For information about...	Refer to page...
<a href="#">Configuration and Image File Management on Your System</a>	3-1
<a href="#">Automated Deployment</a>	3-2
<a href="#">Saving a Configuration</a>	3-2
<a href="#">Executing a Configuration</a>	3-3
<a href="#">Deleting a Configuration Restore-Point or File</a>	3-3
<a href="#">Downloading a File from an FTP, TFTP, or SCP Server</a>	3-4
<a href="#">Downloading a Firmware Image via the Serial Port</a>	3-5
<a href="#">Uploading a Configuration File</a>	3-8
<a href="#">Setting the Boot Firmware Image</a>	3-8
<a href="#">Running a Configuration Script</a>	3-8
<a href="#">Configuration and Image File Display Commands</a>	3-9

### Configuration and Image File Management on Your System

On 7100-Series devices, configuration and image file management includes the following:

- Saving a configuration
- Executing a configuration
- Deleting an image file, configuration file, or script file
- Downloading an image file, configuration file, or a script file
- Uploading a configuration file
- Setting the boot firmware image
- Running a configuration script created on a PC

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

## Automated Deployment

The automated deployment feature allows a newly installed device, with no administrative configuration (default configuration), to automatically obtain the latest firmware revision and configuration from the network. The DHCP client, using a default VLAN and base MAC, obtains a temporary IP address and mask and the IP address of the NetSight server from the DHCP server. Use DHCP Vendor Class Identifier (VCI) code option 60 to identify “Extreme Networks, Inc.” as the vendor. Use DHCP vendor specific code option 125 to specify the information required for the switch to send the trap to start the Netsight OneView application.

Upon obtaining the NetSight IP address, automated deployment sends a notification to NetSight, informing NetSight that its status is READY. NetSight, using its inventory manager, drives the needed changes to the device by determining the configuration to apply and querying the device as to any required upgrades. Refer to the NetSight help page “New Device Configuration in OneView” for NetSight device configuration information.

The automated deployment feature has no CLI input associated with it. Its only dependency is whether the device is currently at the default configuration.

To determine if a device is currently at default configuration, login to the device and a banner will display stating that the system is running with default configuration with a temporary IP address assigned by the DHCP server.

If any CLI configuration is entered for a device configured at default configuration, the device releases the IP address, any SNMP and DHCP settings are backed out, and the device exits automated deployment mode.

To administratively set the device to its default configuration enter the **clear config all** command. If any licenses are configured, also enter the **clear license all** command.

## Saving a Configuration

You can save the 7100-Series device configuration by doing one of the following:

- Creating a configuration restore-point. The configuration restore-point resides on your system. You cannot save a configuration restore-point to a file. Any additional configuration settings that you change after creating this restore-point will not be included when the restore-point configuration is applied, such as when the system reboots. You can configure only one restore-point.

To create a configuration restore-point of the current configuration, use the **set config restore-point** command.

```
set config restore-point <description>
```

- Write the configuration to a file.

To write the configuration to a file, use the **show config outfile** command.

```
show config outfile outfile
```

*outfile* must include the slotN/ local file path directory.

### Example: Creating a Configuration Restore-Point

```
7100 Chassis(rw)-> set config restore-point 25June2009_0800
```

### Example: Creating a Configuration File

```
7100 Chassis(rw)-> show config outfile slot1/newconfig
```

## Executing a Configuration

You can execute the 7100-Series device configuration by doing one of the following:

- Execute the configuration restore-point. Any changes that you made to the configuration after you created the configuration restore-point will be overwritten. See [Procedure 3-1](#).
- Execute a configuration file that was created on, or downloaded to, the 7100-Series device.

### Procedure 3-1 Executing the Configuration Restore-Point

Step	Task	Command(s)
1.	View the index of the configuration restore-point.	<b>show config restore-point</b>
2.	Indicate that the restore-point will be applied when the 7100-Series device reboots. When the 7100-Series device reboots, any configuration changes made after the restore-point was set will be lost.	<b>configure restore-point <i>index</i></b>
3.	Reboot the 7100-Series device.	<b>reset</b>
4.	(Optional) Append the current configuration with the configuration in a previously downloaded or created configuration file.  <b>Note:</b> If you do not specify <b>append</b> , the current running configuration will be replaced with the contents of the configuration file, which will require an automated reset of the chassis.	<b>configure <i>filename</i> [append]</b>

To execute the configuration in a configuration file stored on the 7100-Series device, use the **configure** command.

```
configure filename [append]
```

*filename* must include the slotN/ file path.

The **append** option executes the configuration as an appendage to the current configuration. This is equivalent to typing the contents of the config file directly into the CLI and can be used, for example, to make incremental adjustments to the current configuration.

### Example: Executing a Configuration Restore-Point

```
7100 Chassis(rw)->show config restore-point

Index:          1245935343
Creation Date:  THU JUN 25 13:09:03 2009
Description:    test
7100 Chassis(rw)-> configure restore-point 1245935343
7100 Chassis(rw)->reset
```

### Example: Executing a Configuration File

```
7100 Chassis(rw)->configure slot1/myconfig
```

## Deleting a Configuration Restore-Point or File

You can delete the 7100-Series device configuration by doing one of the following:

- Delete the configuration restore-point. See [Procedure 3-2](#).

- Delete a configuration file.

### Procedure 3-2 Deleting the Configuration Restore-Point

Step	Task	Command(s)
1.	View the index of the configuration restore-point.	<b>show config restore-point</b>
2.	Delete the current restore-point. Because the system currently supports only one restore-point, you must delete the current restore-point before creating a new one.	<b>clear config restore-point <i>index</i></b>
3.	(Optional) Create a new restore-point.	<b>set config restore-point <i>&lt;description&gt;</i></b>

To delete a configuration file, image file, or script file, use the **delete** command.

```
delete filename
```

*filename* must include the slotN/ or images/ file path directory.

#### Example: Deleting a Configuration Restore-Point

```
7100 Chassis(rw)-> clear config restore-point 1245935343
```

#### Example: Deleting a Configuration File

```
7100 Chassis(rw)->delete slot3/myconfig
```

#### Example: Deleting an Image File

```
7100 Chassis(rw)->delete images/010300
```

## Downloading a File from an FTP, TFTP, or SCP Server

You can download an image file, a configuration file, or a script file from an FTP, TFTP, or SCP server to the 7100-Series device.

To download an image file, configuration file, or script file from an FTP, TFTP, or SCP server, use the **copy** command.

```
copy source destination
```

– *source* is the URL of an FTP, TFTP, or SCP server.

– *destination* is the local file path. For a configuration or script file, *destination* must include slotN/.

The 7100-Series module to which a configuration file is downloaded must have the same hardware configuration as the 7100-Series module from which it was uploaded.

For reasons of security, passwords are not allowed in **copy** command URLs. A password prompt displays upon entering a **copy** command. For example:

```
7100 Chassis(rw)->copy scp://doc@banshee.extremenetworks.com:22/myconfig
slot3/myconfig
```

```
Password:
```

```
#####
```

```
7100 Chassis(rw)->
```

Once you have downloaded an image file, set the device to load the new image file at startup using the **set boot system** command. See “[Setting the Boot Firmware Image](#)” on page 3-8.

For information on downloading

### Example: Downloading an Image File

```
7100 Chassis(rw)->copy tftp://134.141.89.34/ets-mtxe7-msi newimage
```

### Example: Downloading a Configuration File

```
7100 Chassis(rw)->copy tftp://134.141.89.34/myconfig slot3/myconfig
```

## Downloading a Firmware Image via the Serial Port

Besides using FTP, TFTP, or SCP for downloading firmware images, you can also download firmware images via the serial (console) port. This procedure is an out-of-band operation that copies the firmware through the serial port to the device. It should be used in cases when you cannot connect to the device to perform the in-band **copy** download procedure via FTP, TFTP or SCP. Serial console download has been successfully tested with the following applications:

- HyperTerminal
- TeraTerm

Any other terminal applications may work but are not explicitly supported.

---

### Important Notice

The 7100-Series device allows you to download and store multiple image files. This feature is useful for reverting back to a previous version in the event that a firmware upgrade fails to boot successfully. After downloading firmware as described above, you can select which image file you want the device to load at startup using the **setboot** command in the System Image Loader menu or the **set boot system** command.

---

To download device firmware via the serial (console) port, proceed as follows:

1. With the console port connected, power up the device. The following message displays:

```
Boot ROM Initialization, Version 01.00.02

Copyright (c) 2003 Enterasys Networks, Inc.
SDRAM size: 1024 MB
Testing SDRAM... PASSED.
Loading Boot Image: 01.00.19... DONE.
Uncompressing Boot Image... DONE.
```

2. Once the boot image is finished uncompressing, you receive a message indicating you have 3 seconds to access the bootloader menu by pressing any key. Press a key and the system image loader prompt displays:

```
###You have 3 seconds to access the bootloader menu###
Press any key to enter System Image Loader menu
PressAnyKey
[System Image Loader]:
```

3. To display help for all the system image loader mode commands, enter a question mark (?):

```
[System Image Loader]:?
?, help - print this list
boot - boot (load and go)
delete - delete an image file
download - start ZMODEM download
list - display available images
```



```

log                - message log
setbaud <rate>    - set baud rate, (9600,38400,57600,115200)
setboot <filename> - change boot image file
showboot          - display boot image file
clearnvram        - clear persistent storage

```

[System Image Loader]:

4. Use the **list** command to display the images currently on this device.
5. The baud rate can be set to 9600, 38400, 57600, or 115200. Using the **setbaud** command, set the baud rate to **115200**:

```

[System Image Loader]: setbaud 115200
###Change the baud of the terminal program to 1152000###
[System Image Loader]:

```

6. Use the **download** command to start the ZMODEM receive process. Send the image file using the ZModem protocol from your terminal application. (This procedure will vary depending on your application.) When the ZModem download is finished, the following message displays:

```

[System Image Loader]: download
Preparing to receive file...
**XXXXXXXXXXXXXXXXXXXX
###Start the ZMODEM transfer from the terminal software###
Writing file...
Download successful.
[System Image Loader]:

```

7. Use the **list** command to confirm the images that are currently on the device, and confirm the image currently listed as the boot image. If the current boot image is not the image you want to boot with, use the **setboot filename** command to set the correct boot image:

```

[System Image Loader]: list
Filename:      720010001 (Boot)
Version:      07.20.01.0001
Size:         4527490 (bytes)
Date:         FRI DEC 10 15:32:24 2010
Checksum:     d89ace409317bc765789fce1c73b8745
Compatibility: listOfCompatibleDevices

Filename:      720010025
Version:      07.20.01.0025
Size:         4529790 (bytes)
Date:         THU DEC 09 22:38:54 2010
Checksum:     6ccaaf8a5b77d7d34c6c3d972b381024
Compatibility: listOfCompatibleDevices
[System Image Loader]:setboot 720010025
[System Image Loader]:list
Filename:      720010001
Version:      07.20.01.0001
Size:         4527490 (bytes)

```

```
Date:          FRI DEC 10 15:32:24 2010
Checksum:     d89ace409317bc765789fce1c73b8745
Compatibility: listOfCompatibleDevices
```

```
Filename:     720010025 (Boot)
Version:     07.20.01.0025
Size:        4529790 (bytes)
Date:        THU DEC 09 22:38:54 2010
Checksum:    6ccaaf8a5b77d7d34c6c3d972b381024
Compatibility: listOfCompatibleDevices
[System Image Loader]:
```

- When a device is booted, the device baud rate is reset to 9600. Reset the terminal application baud rate to 9600 so that it will continue to display output from the device:

```
[System Image Loader]: setbaud 9600
[System Image Loader]:
```

- Use the **boot** command to boot the image:

```
[System Image Loader]: boot
###The unit will boot normally###
```

/flash0/ - Volume is OK

```
Loading 61205...          DONE.
Uncompressing System Image...  DONE.
Loading System Image...   DONE.
Initializing Platform Hardware
```

.  
.
.  
.

```
Enterasys Networks, Inc.
50 Minuteman Rd.
Andover, MA 01810-1008 USA
```

```
Phone:  +1 978 684 1000
E-mail:  support@enterasys.com
WWW:    http://www.enterasys.com
(c) Copyright Enterasys Networks, Inc. 2014
Chassis Serial Number:    00e063937c7d
Chassis Firmware Revision: 07.20.01.0025
Username:
```



**Note:** If you reboot without specifying the image to boot with **setboot** as described above, the device will attempt to load whatever image is currently stored in the bootstring via the **setboot system** command. If the device cannot find the image, or it is not set, it will search through available images and attempt to boot the newest one. If the device finds and successfully boots an image file, it will set the bootstring to the name of that image file.

## Uploading a Configuration File

You can upload a configuration file from the 7100-Series device.

To upload a configuration file, use the **copy** command.

```
copy source destination
```

- *source* is the local file path and must include slotN/.
- *destination* is the URL of an FTP, TFTP, or SCP server.

### Example

```
7100 Chassis(rw)->copy slot3/myconfig ftp://134.141.89.34/myconfig
```

## Setting the Boot Firmware Image

You can set the boot firmware image, which is the image that will be loaded automatically after the system has been reset.

To set the boot firmware image, use the **set boot system** command.

```
set boot system filename
```

The system must be reset by software for the new boot image to take effect at startup. If the chassis is powered OFF and then back ON, the current active image will just reload at startup.

Although it is not necessary to choose to reset the system and activate the new boot image immediately, the CLI will prompt you whether or not you want to do so. You can choose “Yes” at the question prompt to have the system reset and load the new boot image immediately, or choose “No” to load the new boot image at a later scheduled time by issuing one of the following commands: **clear config**, **reset**, or **configure**. The new boot setting will be remembered through resets and power downs, and will not take effect until the **clear config**, **reset**, or **configure** command is given.

### Example

```
7100 Chassis(rw)->set boot system newimage
```

This command can optionally reset the system to boot the new image.

```
Do you want to reset now (y/n) [n]?y
```

```
Resetting system ...
```

## Running a Configuration Script

You can run a configuration script that you have downloaded to the 7100-Series device. See [Procedure 3-3](#).

### Procedure 3-3 Running a Configuration Script

Step	Task	Command(s)
1.	Download the configuration script. <i>source</i> is the URL of an FTP, TFTP, or SCP server. <i>destination</i> is the local file path and must include slotN/.	<b>copy</b> <i>source destination</i>
2.	Run the configuration script.	<b>script</b> <i>filename</i> [ <i>arg1</i> ] [ <i>arg2</i> ] [ <i>arg3</i> ] [ <i>arg4</i> ] [ <i>arg5</i> ] [ <i>arg6</i> ] [ <i>arg7</i> ]

## Example

This example uses the **copy** command to copy the script file named “setport.scr” from IP address 10.1.221.3 to slot 4. Next, the contents of the file is displayed with the **show file** command. The script file requires two arguments, a port string (%1) and a VLAN id (%2). Finally, the script is executed, by specifying ge.1.1 as the first argument and 100 as the second argument.

```
7100 Chassis(rw)->copy tftp://10.1.221.3/setport.scr slot4/setport.scr
```

```
7100 Chassis(rw)->show file slot4/setport.scr
```

```
set port alias %1 script_set_port
```

```
set port vlan %1 %2 modify-egress
```

```
set port jumbo enable %1
```

```
set port disable %1
```

```
set port lacp port %1 disable
```

```
7100 Chassis(rw)->script slot4/setport.scr ge.1.1 100
```

## Configuration and Image File Display Commands

Table 3-1 lists configuration and image file display commands for 7100-Series devices.

**Table 3-1 Configuration and Image File Management and Display Commands**

Task	Command
Display the index, creation date, and description of the currently configured restore-point. If “(Boot)” is listed after the index entry, this restore-point will be used when the system reboots next.	<b>show config restore-point</b>
Display the firmware image the system will load at the next system reset.	<b>show boot system</b>
List files stored in the file system.	<b>dir [filename]</b>
Display the contents of an image or configuration file.	<b>show file filename</b>
Display the system configuration.	<b>show config [all] [facility]</b>

Refer to the device’s *CLI Reference Guide* for a description of the output of each command.



## High Availability Firmware Upgrade (HAU) Configuration

This chapter provides information about configuring and monitoring a High Availability Firmware Upgrade (HAU) on 7100-Series devices.

For information about...	Refer to page...
<a href="#">Using High Availability Firmware Upgrade in Your Network</a>	4-1
<a href="#">Implementing HAU</a>	4-2
<a href="#">High Availability Upgrade Preconditions</a>	4-2
<a href="#">System Limitations During a High Availability Upgrade</a>	4-3
<a href="#">HAU Configuration Overview</a>	4-3
<a href="#">Terms and Definitions</a>	5-13

### Using High Availability Firmware Upgrade in Your Network



**Note:** The 7100-Series HAU feature support is limited to a Virtual Switch Bond (VSB) configuration. Refer to [Chapter 5, Virtual Switch Bonding \(VSB\) Configuration](#) for VSB configuration details.

High Availability Firmware Upgrade (HAU) is an 7100-Series feature that provides for a rolling firmware upgrade for maintenance releases that are HAU compatible with the current system firmware.

There are two methods for loading a system firmware image:

- Standard – The specified image is loaded after a system reset
- High Availability – Provides a rolling firmware upgrade

Using the standard upgrade method, the image is loaded automatically after the system has been reset. Using the HAU method, each 7100-Series chassis in the VSB configuration is assigned to a unique HAU group. The chassis not belonging to the HAU group currently being upgraded remains operational. After the first chassis completes its upgrade, the second chassis in the VSB configuration begins its upgrade. Whether the upgrade procedure uses the standard or HAU method, during the upgrade, one of the 7100-Series chassis is not operational. To avoid potential feature conflicts between multiple firmware versions, the HAU firmware upgrade feature is limited to maintenance firmware upgrades and will not be available when upgrading to major feature releases.

A number of switch applications require a period of time after the completion of a firmware upgrade to become fully operational. During this period, some resources associated with the just

completed HAU group will not be available. By default a 5 second delay takes place between the completion of an HAU group upgrade and the reset of the next HAU group.

When specifying the image to use when booting the system, you can optionally specify whether a standard or high availability system upgrade should be performed.

If no system boot method is specified when setting the system boot image, there are three configurable HAU default modes that determine upgrade behavior:

- A high availability upgrade will never be performed. In this case the standard upgrade method is used in all cases. This is the default system boot behavior.
- A high availability upgrade will be performed if all HAU preconditions are met. If any HAU preconditions are not met, a standard upgrade is performed. See [“High Availability Upgrade Preconditions”](#) on page 4-2 for HAU precondition details.
- A high availability upgrade will be performed if all HAU preconditions are met. If any HAU preconditions are not met, no upgrade is performed. See [“High Availability Upgrade Preconditions”](#) on page 4-2 for HAU precondition details.

You can not halt a high availability upgrade once it has begun, but you can force all remaining HAU groups that have not yet been upgraded to immediately upgrade. See [“Forcing Early Completion of a Running HAU”](#) on page 4-5 for details.

## Implementing HAU

To implement HAU:

1. Optionally, modify the default HAU mode. By default a standard (non-high availability) upgrade is performed, unless high availability is set when specifying the boot image.
2. Assure that each chassis in the VSB configuration is assigned to its own HAU group.
3. Optionally, modify the delay between the completion of one HAU group upgrade and the start of the next HAU group upgrade.
4. Specify the system image that will boot the next time the system resets, and optionally specify whether the image will load using the standard or high availability method. If you do not specify the boot mode, the image will load based upon the default HAU mode referred to in Step 1.

## High Availability Upgrade Preconditions

The following preconditions must be met for an high availability upgrade to occur:

- **HAU Compatibility Key** - The target image must have the same HAU Compatibility Key as the active image. To display the HAU key, use the **dir** command, specifying the image to display, or use the **dir** command **image** option to display all images. The HAU key field in the display specifies whether the image displayed is compatible with the current image. If “HAU compatible” is appended to the key field, a high availability upgrade can be performed between the displayed image and the current image.
- **Configuration restore-points** - Configuration restore-points may be set, but must not be configured. A configured restore-point would cause upgraded slots to boot with different configuration data, and all slots must be running the same configuration data.
- The 7100-Series chassis are members of a Virtual Switch Bonded (VSB) system with each chassis configured for a unique HAU group.

## System Limitations During a High Availability Upgrade

Changes to system configuration cannot be performed while a high availability upgrade is in progress. While a high availability upgrade is running:

- All SNMP set operations are rejected. A “noAccess” reason will be given for the rejection.
- All CLI commands are unavailable with the exception of:
  - reset
  - loop
  - show
  - exit
  - dir
  - history
  - ping
  - traceroute
  - telnet
  - ssh
  - set boot high-availability force-complete

## HAU Configuration Overview

For information about...	Refer to page...
<a href="#">Configuring System Boot Image and Mode</a>	4-3
<a href="#">Configuring HAU Default Mode</a>	4-4
<a href="#">Configuring HAU Groups</a>	4-4
<a href="#">Configuring a Delay Between HAU Group Upgrades</a>	4-5
<a href="#">Disabling a Configured HAU</a>	4-5
<a href="#">Forcing Early Completion of a Running HAU</a>	4-5
<a href="#">High Availability Firmware Upgrade in a Virtual Switch Bonded System</a>	4-6

## Configuring System Boot Image and Mode

When a system is powered on or reset, the current system boot image is loaded on to all system modules. To perform a system upgrade, change the current system boot image to the upgrade image, also referred to as the target image. Image upgrade can occur immediately, the next time the system boots, or by issuing a **reset** command. When specifying the new target image, you can optionally, specify the system boot mode parameter:

- **Standard** – All system slots are simultaneously upgraded taking the system out of operation for the duration of the upgrade. This is a non-high availability upgrade.
- **High-availability** – Providing all HAU preconditions are met, HAU groups are upgraded sequentially (one after another, but not in any specific order). If any HAU precondition is not met, an upgrade does not occur. See “[High Availability Upgrade Preconditions](#)” on page 4-2 for HAU precondition details.



If the system boot mode is not specified, the boot mode is determined by the HAU default mode configuration. By default, the HAU default mode executes a **standard** system upgrade. See “[Configuring HAU Default Mode](#)” on page 4-4 for HAU default mode details.

Use the **set boot system** *image-file-name* command in any command mode to set the target image for the firmware upgrade, optionally specifying the system boot mode.

Use the **set boot system** *image-file-name* **high-availability** command in any command mode to set the target image for the firmware upgrade and enable HAU, assuring that if a high availability upgrade is possible it will be performed, otherwise no upgrade will occur.

When entering the **set boot system** command, you are asked if you want to reset the system or start a high availability upgrade depending upon the command variation you entered. Answering **y** will immediately reset the system and begin the upgrade. Answering **n** will perform the upgrade the next time the system is reset.

Use the **show boot system** command to determine the current boot image.

Use the **dir** *image-file-name* command to display image details for the specified image.

Use the **dir images** command to display details for all available images on the system.

## Configuring HAU Default Mode

HAU default mode determines HAU behavior if a system boot mode is not set when configuring the system boot image. See “[Configuring System Boot Image and Mode](#)” on page 4-3 for system boot mode details. There are three HAU default modes:

- **never** – A standard (non-high availability) upgrade is always performed unless over-ridden by the system boot mode **high-availability** setting
- **if-possible** – A high availability upgrade is always performed unless:
  - All HAU preconditions are not met, in which case a standard upgrade is performed
  - Over-ridden by the system boot mode **standard** or **high-availability** settings
- **always** – A high availability upgrade is always performed unless:
  - All HAU preconditions are not met, in which case no upgrade occurs
  - Over-ridden by the system boot mode **standard** setting

See “[High Availability Upgrade Preconditions](#)” on page 4-2 for HAU precondition details.



**Note:** HAU default mode should always be set to **never** unless you intend to perform a high availability upgrade. An **if-possible** or **always** HAU default mode setting in conjunction with no system boot mode specified results in a high availability firmware upgrade each time you reboot your system, if all HAU preconditions are met.

If you want an HAU default mode change to affect a firmware upgrade, the change must take place before configuring a pending upgrade. Changing the HAU default mode after setting the system boot configuration (using the **set boot system** command) has no affect on a pending firmware upgrade.

Use the **set boot high-availability default-mode** command in any command mode to set the HAU default mode.

## Configuring HAU Groups

The HAU group feature determines which chassis will be simultaneously upgraded. All chassis within the same HAU group are simultaneously upgraded. Each system slot belongs to an

HAU group. HAU occurs one HAU group at a time. By default, there is one slot per group. Therefore, the default HAU behavior is to upgrade each system slot one at a time.

All essential system capabilities on the device should be configured across multiple groups. For example, all LAGs configured on the device should provide sufficient redundancy between HAU groups for packets to continue forwarding on the LAG using chassis belonging to HAU groups that are not upgrading.

Use the **set boot high-availability group** command in any command mode to configure an HAU group, specifying the group ID and the system slots that will be members of the HAU group.

## Configuring a Delay Between HAU Group Upgrades

When the firmware upgrade of an HAU group completes, depending upon the applications that are configured on the module, it is possible for the next HAU group to begin a firmware upgrade prior to protocols or applications on the just completed HAU module becoming fully operational. Under normal operation there is an approximately 5 second delay between the completion of one HAU group upgrade and the start of the next group upgrade. You can configure an additional delay of up to 600 seconds between the upgrade completion of one HAU group and the beginning of a high availability upgrade for the next HAU group.

Use the **set boot high-availability delay** command in any command mode to set a delay in seconds between the upgrade completion of any HAU group and the beginning of the next HAU group upgrade.

## Disabling a Configured HAU

You can disable a pending high availability upgrade by:

- Setting the boot image back to the active image using the **set boot system *active-image*** command
- Deleting the boot image using the **delete *target-image*** command
- Converting the pending high availability upgrade to a standard upgrade by re-issuing the boot command, specifying the target image and the standard system boot mode

After performing one of the methods for disabling an HAU configuration, verify that the HAU status is disabled by using the **show boot high-availability** command.

You cannot disable a high availability upgrade that is running. You can however force the simultaneous upgrade of all remaining non-upgraded HAU groups. See [“Forcing Early Completion of a Running HAU”](#) on page 4-5 for details.

## Forcing Early Completion of a Running HAU

You cannot disable a high availability upgrade or revert an image back to the original system image on a high availability upgrade that is running. You can however accelerate the upgrade process, by forcing the simultaneous upgrade of all remaining non-upgraded HAU groups.

This should not be considered a normal HAU procedure. It should be assumed that forcing the early completion of HAU will degrade the operational capabilities of the system depending upon the system resources taken out of service.

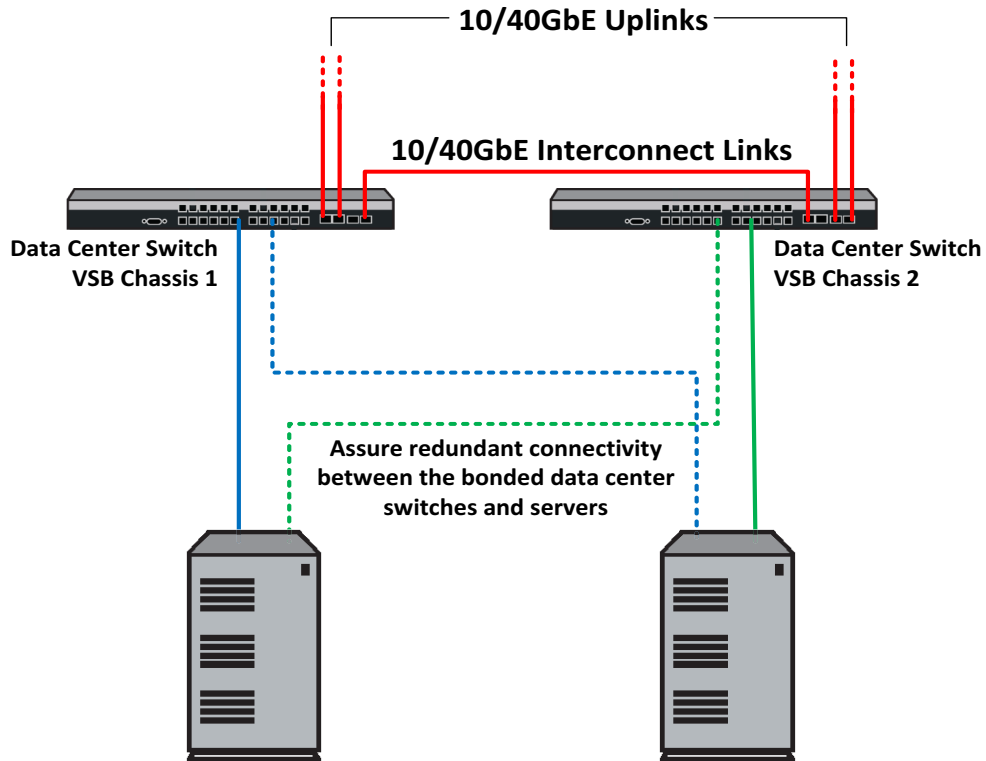
Use the **set boot high-availability force-complete** command in any command mode to force the simultaneous upgrade of all non-upgraded HAU groups in the system.

## High Availability Firmware Upgrade in a Virtual Switch Bonded System

HAU in a Virtual Switch Bonding context assumes that the VSB system is properly configured and fully operational. See “[Virtual Switch Bonding \(VSB\) Configuration](#)” on page 5-1 for VSB configuration details.

Figure 4-1 on page 4-6 presents a 7100-Series HAU VSB configuration in a data center setting. When connecting the standalone bonded data center switches to data center servers, assure there is connectivity from each 7100-Series switch to each connected server for all configured LAGs. When the 7100-Series device is taken down for the upgrade, all connectivity from that device is taken out of service for the duration of the upgrade. If redundant connectivity is not present between the servers and the operational bonded switch, any service associated with the missing connectivity will not be available to one or more servers.

**Figure 4-1 VSB System High Availability Firmware Upgrade**



## Configuring HAU

This section provides a table of HAU default values and a procedure for configuring an HAU system.

Table 4-1 lists HAU default values.

**Table 4-1 Default HAU Parameters**

Parameter	Description	Default Value
system boot mode	Determines firmware load or upgrade behavior when the system is powered on or reset.	standard

**Table 4-1 Default HAU Parameters (continued)**

Parameter	Description	Default Value
HAU delay	A configurable delay in seconds between the upgrade completion of one HAU group and the beginning of the next HAU group upgrade.	0 seconds
HAU default mode	The boot mode that is used when no boot mode is specified when setting the system boot image.	never

[Procedure 4-1](#) describes HAU configuration on the Extreme Networks 7100-Series devices. All commands used to configure HAU can be entered in any command mode.

#### Procedure 4-1 Configuring HAU

Step	Task	Command(s)
1.	Optionally, modify the default HAU mode. By default a standard (non-high availability) upgrade is performed, unless high availability is set when specifying the boot image.	<b>set boot high-availability default-mode</b> {never   if-possible   always}
2.	Optionally, modify the default HAU group configuration by placing multiple slots into some or all HAU groups.	<b>set boot high-availability group</b> <i>group-id</i> <i>slot(s)</i>
3.	Optionally, set a delay between the end of an HAU group upgrade and the beginning of the next HAU group upgrade.	<b>set boot high-availability delay</b> <i>delay</i>
4.	Specify the system image that will boot the next time the system resets, and optionally specify whether the image will load using the standard or high availability method.	<b>set boot system</b> <i>filename</i> [standard   high-availability]

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

The following is an example of the HAU aspect of a high availability upgrade configuration, based upon the system setup presented in [Figure 4-1](#) on page 4-6. This configuration example:

- Assumes that VSB and LAGs have been correctly configured
- Leaves the high availability default mode unchanged because the high availability method was specified in the **set boot system** command
- Leaves the default HAU group configuration of each chassis configured in its own HAU group unchanged
- Sets the delay between the upgrade completion of one HAU group and the beginning of another to 15 seconds on each chassis
- Sets the boot system mode to high-availability for image **7100-79101-0003** on each chassis

```
7100-Series(rw)->set boot high-availability delay 15
```

```
7100-Series(rw)->set boot system 7100-79101-0003 high-availability
```

## Terms and Definitions

Table 4-2 lists terms and definitions used in this security mode configuration discussion.

**Table 4-2 HAU Configuration Terms and Definitions**

Term	Definition
active or current image	The image that is currently running or configured to load on the system
HAU compatibility key	A key associated with each image on the system that must be the same for both the target and active images for a high availability upgrade to take place.
HAU default mode	The boot mode that is used when no boot mode is specified when setting the system boot image.
HAU delay	A configurable delay in seconds between the upgrade completion of one HAU group and the beginning of the next HAU group upgrade.
HAU force complete	A high availability upgrade feature that forces all HAU groups that have not yet upgraded to immediately and simultaneously upgrade.
HAU group	A configuration of one or more slots into a group that results in group members being simultaneously upgraded during a high availability upgrade, leaving slots that are members of all non-upgrading groups operational.
high availability boot mode	A system boot mode option that upgrades HAU groups one at a time, if all HAU preconditions are met.
High Availability Firmware Upgrade (HAU)	An Extreme Networks feature that allows a system to retain high availability while being upgraded to an HAU compatible release. By contrast, all system connectivity and features are temporarily lost during a standard upgrade.
interconnect port	10GbE port that plays the same role as the backplane fabric in a non-VSB chassis by providing distribution between VSB system chassis.
standard boot mode	A system boot mode option that simultaneously upgrades all system slots, taking the system out of operation for the duration of the upgrade.
system boot mode	Determines firmware load or upgrade behavior when the system is powered on or reset.
target image	The image that the system will upgrade to on the next system reset.
Virtual Switch Bonding (VSB)	An Extreme Networks S-Series feature that aggregates two like chassis into a single virtual network device.
VSB chassis	One of two chassis configured for VSB that make up a VSB system.
VSB system	The aggregation of two chassis configured for VSB and connected by one or more 10GbE interconnect ports.

## Virtual Switch Bonding (VSB) Configuration

This chapter provides information about configuring and monitoring Virtual Switch Bonding (VSB) on 7100-Series devices.

For information about...	Refer to page...
<a href="#">Using Virtual Switch Bonding in Your Network</a>	5-1
<a href="#">Implementing VSB Stacking</a>	5-4
<a href="#">VSB Configuration Overview</a>	5-4
<a href="#">Configuring VSB</a>	5-11
<a href="#">Terms and Definitions</a>	5-12

### Using Virtual Switch Bonding in Your Network

Virtual Switch Bonding (VSB) stacking is an Extreme Networks 7100-Series feature that allows for the aggregation of links on up to 8 physical chassis, providing redundancy, while at the same time allowing ports on all chassis to pass data concurrently. VSB aggregates multiple chassis into a single virtual network device. VSB joins multiple chassis into a single system by extending each chassis' distribution to the next chassis in the stack using one or more 40GbE port interconnect links. Any mix of 7100-Series family units – 7100G and 7100K can be VSB stacked.

The 40GbE ports connecting the physical chassis are designated as bonding ports on each chassis and create the virtual backplane that ties the two physical chassis together.

VSB stacking supports the Link Failure Response (LFR) feature. LFR is a mechanism designed to protect against what is often referred to as the “split brain” problem using LFR priority configuration as a means of determining which surviving stack segment will remain up when a link or chassis failure occurs. This problem occurs when a VSB stack splits into two or more segments, due to the failure of a bonding link or individual chassis, resulting in duplication of managements network interfaces. This scenario can result in serious disruptions to the customer network.

40GbE VSB stacking configured ports are always set as interconnect ports. 1 or 10 GbE VSB configured ports are always set as LFR monitor ports.




**Note:** A closed ring configuration configures an interconnect link between the first and last 7100-Series chassis in the stack. Closed ring configurations are not required, but are highly recommended. A system or interconnect link failure in an open loop configuration can result in separate bonded systems both using the same IP address. When a system or interconnect link failure occurs in a closed ring configuration, this issue will not occur.

VSB is typically used in a data center between switches LAGed to a server on one side and to network devices on the other.

Figure 5-1 presents an overview of a four chassis VSB stacking system in a data center context. The top two 7100K-Series chassis (Chassis 1 and 2) have four 40GbE links, two of which are used as redundant VSB stacking interconnect links. Chassis are configured in order from top to bottom as VSB chassis 1, 2, 3, and 4. These chassis are members of VSB system 1. Between chassis 1 and 2, two 40GbE interconnect links are used providing a redundant interconnect link. Chassis 1 uses a third 40GbE port to provide a closed ring configuration between chassis 1 and 4. There is a single VSB interconnect link between chassis 2-to-3, 3-to-4, and 4-to-1. The fourth 40GbE port on chassis 1 and 2 are used as uplinks to the network. 10GbE ports on chassis 3 and 4 can also be used uplink ports to the network.

In Figure 5-1, a server with two NICs installed achieves redundancy through a direct connection, distributing the two NIC connections between the two available VSB system slots. A single NIC server achieves redundancy through an intermediate switch that distributes multiple connections between the two VSB system slots.

**Figure 5-1 VSB Stacking Data center Configuration Overview**



Assuming **fg.x.1** of each chassis is connected to **fg.x.2** of the next chassis, with **fg.4.1** creating a closed ring by interconnecting to **fg.1.2**, the show bonding display for Figure 5-1 would be:

```
7100 Chassis(rw) ->show bonding
```

```

Global Bonding State      : enabled
Max Bonded Chassis      : 8
Max Bonded System Slot   : 8
Slots active in System   : 1-8
Bonding Admin System MAC: 00-00-71-00-00-71
Bonding Oper System MAC  : 00-00-71-00-00-71
Link Failure Response    : disabled

```

System		Shared	LFR	Active Slot IDs
Chassis Identifier	Serial Number	Secret	Priority	For This Chassis
1	999 TOR000035	no	10	1
2	999 124700016845	no	20	2
3	999 13190048685E	no	30	3
4	999 13170490685E	no	40	4

Port	Admin Status	Partner Port	Oper Mode	Oper Status
fg.1.1	enabled	fg.2.2	bonding	up
fg.1.2	enabled	fg.4.1	bonding	up
fg.2.1	enabled	fg.3.2	bonding	up
fg.2.2	enabled	fg.1.1	bonding	up
fg.3.1	enabled	fg.4.2	bonding	up
fg.3.2	enabled	fg.2.1	bonding	up
fg.4.1	enabled	fg.1.2	bonding	up

```
7100 Chassis (rw) ->
```



**Note:** A closed ring interconnect is not required, but if you do not close the ring and an interconnect or a system failure occurs, the remaining systems could be divided, causing two systems to reside in your network with the same IP address.

A VSB stacking system is managed by a single IP address and behaves as if it is one chassis with multiple slots. Once globally enabled, VSB system IP address configuration is the same as for a non-VSB multi-slot system.

The VSB stacking system provides gains in network resiliency, performance, and management through:

- A greater number of ports in a single system
- Location redundancy
- Added bandwidth

Though bandwidth and port capacities scale in a VSB system, feature capacities such as route, MAC address tables and user capacities remain the same as in a single chassis system.

When activating an advanced routing license on a 7100-Series VSB stacking system, you must activate the license on all chassis in the system.



Using the High Availability feature, slots can be grouped such that forwarding is maintained while individual slot groups are firmware upgraded, so long as a group that is not being upgraded continues to provide the interconnect between VSB chassis.

A non-default MAC address can be manually set prior to globally enabling the VSB system. VSB stacking chassis can be configured with unique MAC addresses or the same MAC address. It is recommended that you configure all chassis in the VSB system with the same MAC address.

## Implementing VSB Stacking

To implement VSB:

1. Select the chassis that will be the physical chassis members of the VSB stacking system.
2. Assure that all chassis intended for the VSB stacking system are configured for the same firmware image. Different firmware images in the stack will result in multiple reboots for the affected chassis and may result in a host master election of a chassis with a non-desired firmware version.
3. Configure each chassis' VSB chassis ID and the ID of the VSB system the chassis belong to.
  - If LFR is being used, optionally, change the LFR priority for one or more chassis in the stack (the LFR priority setting must be unique for each chassis)
4. Identify the 40GbE ports that will be used to interconnect the chassis. One VSB connection is required between each chassis; multiple VSB interconnections can be used between 7100K-Series chassis due to the availability of four 40GbE ports. VSB interconnect ports should be selected taking into consideration the optimization of bandwidth usage and redundancy. Interconnect the first and last chassis in the system to form a closed ring interconnect configuration.
5. Enable bonding on the selected 40GbE interconnect ports (**fg.1.x** where *x* is the port number).
6. If you are using LFR, enable bonding on the 1 or 10GbE monitor ports.
7. Optionally change the VSB MAC address from the default value. The VSB MAC address defaults to an internal MAC address associated with chassis 1. It is recommended that all chassis in the VSB stacking system have the same MAC address.
8. If 1 or 10GbE monitor ports were enabled, enable LFR on the system.
9. After completing steps 1 – 8, globally enable the VSB system. The system is reset. If High Availability firmware upgrade is configured prior to enabling VSB, the HAU configuration is lost.
10. Optionally configure the VSB system for High Availability firmware upgrades.
11. When you have completed port level configuration on each chassis, make sure you archive a copy of the configuration should you need to replace a failed unit.

## VSB Configuration Overview

For information about...	Refer to page...
<a href="#">VSB Chassis Configuration</a>	5-5
<a href="#">VSB Interconnect Link Configuration</a>	5-5
<a href="#">Link Failure Response (LFR) Configuration</a>	5-6
<a href="#">VSB System MAC address Configuration</a>	5-8

For information about...	Refer to page...
<a href="#">Licensing</a>	5-9
<a href="#">Globally Enabling and Disabling the VSB System</a>	5-9
<a href="#">High Availability Firmware Upgrade</a>	5-9
<a href="#">Applying a VSB Configuration File to a Replacement VSB Chassis</a>	5-10

## VSB Chassis Configuration

VSB chassis configuration must take place before the system is VSB globally enabled. Prior to a VSB system being globally enabled, you can modify the VSB chassis configuration as needed. Once a VSB system is globally enabled using the **set bonding** command (See [“Globally Enabling and Disabling the VSB System”](#) on page 5-9), VSB specific chassis configuration can not be modified. VSB must be disabled globally, reverting each chassis to a non-VSB state, in order to make modifications to the VSB chassis configuration.

There are two values associated with a VSB chassis configuration:

- Chassis ID – Identifies the physical chassis member of the VSB system.
- System ID – Identifies the VSB system made up of between two and eight physical chassis

VSB chassis ID valid values are **1 - 8**. The VSB chassis ID both identifies the VSB chassis and determines chassis slot assignments within a globally enabled VSB system. After the chassis ID and system ID have been configured, but prior to globally enabling VSB, slot numbering does not differ from standard slot numbering (all 7100-Series slots are labeled slot 1). Once the system has been globally enabled for VSB, the chassis 1 slot is numbered 1, the chassis 2 slot is numbered 2, and so on for all remaining chassis.

The VSB system ID is configured with the same value on each chassis in a given VSB system. The VSB system ID is used to identify the bonded system as a whole. VSB does not enforce unique VSB system ID values between systems, but it is highly recommended that each system be configured with a unique value for management purposes. Valid VSB system ID values are **1 – 18446744073709551615**.

A VSB secret can be configured to add security to the VSB links. The VSB secret can be up to 32 printable characters long. If a space is used, the secret must be enclosed in double quotes (“”). The secret can be created or modified at any time. The secret can be overwritten without first clearing it. If a secret only exists on a single chassis, the VSB system will segment until the same secret is present on both physical chassis in the system. A VSB secret is configured using the **secret** option of the **set bonding chassis** command.

Use the **set bonding chassis** command to configure a VSB chassis.

Use the **clear bonding chassis** command to clear the VSB chassis configuration, only if VSB is not globally enabled on the chassis.

## VSB Interconnect Link Configuration

In a VSB enabled system, 40GbE port interconnections provide distribution between chassis. VSB supports either 2 (7100G-Series) or 4 (7100K-Series) 40GbE interconnection links between chassis. It is recommended that you consider optimization of interconnect redundancy and optimization of bandwidth when determining the number and location of VSB interconnect links configured between bonded chassis.

Configure the VSB chassis using the **set bonding chassis** command (see [“VSB Chassis Configuration”](#) on page 5-5 for details) before enabling VSB 40GbE interconnect ports.

Use the **set bonding port enable** command, specifying the 40GbE interconnect port, to enable the VSB interconnect port. The VSB interconnect port must be enabled on all chassis in the VSB stacking system. Port designation is based upon the standard physical chassis slot and port designation until VSB is globally enabled for the VSB system (See “[Globally Enabling and Disabling the VSB System](#)” on page 5-9). A minimum of one VSB interconnect link must be configured between bonded chassis before globally enabling VSB for the intended system to bond. Once VSB is globally enabled, VSB interconnect links can be added or modified.

See “[VSB Chassis Configuration](#)” on page 5-5 for details concerning chassis slot numbering before and after a VSB system is globally enabled.

In a VSB system of  $x$  chassis that is globally enabled:

- Chassis 1 ports are specified as **fg.1.x** where  $x$  is the port number
- Chassis 2 ports are specified as **fg.2.x** where  $x$  is the port number
- Chassis  $x$  ports are specified as **fg.x.y** where  $x$  is the chassis number and  $y$  is the port number

When modifying interconnect ports in a globally enabled VSB system, use the globally enabled port designation to specify ports.

## Link Failure Response (LFR) Configuration

The LFR feature determines which chassis stack segment front-panel ports will remain enabled should all VSB interconnect links between the VSB chassis go down. All chassis in an operational VSB system stack use the same IP address and function as a single system with the 40GbE interconnect links acting as a virtual backplane for the system. Should all VSB interconnect links between two chassis go down and multiple stack segments remain operational, the stack segments would function as independent network devices with the same IP address.

In a closed ring interconnect configuration as shown in [Figure 5-1](#) on page 5-2, you would need two points of failure for the stack to form multiple segments. The interconnect between the first and last chassis would maintain a single contiguous system in a single point of failure scenario. In a non-closed ring interconnect configuration, any point of failure other than for the first or the last chassis in the system would create multiple stack segments.

When configuring a system of two 7100-Series devices in a VSB system, the LFR feature allows 1 or 10GbE ports to be designated as VSB monitor links that operate in a standby mode to the primary 40GbE VSB ports. The VSB monitor link provides dedicated redundant control plane connectivity and is used only as a backup communication path between bonded chassis in the unlikely event that all of the primary VSB interconnect links fail or become unavailable. When the primary 40GbE VSB ports are down, the VSB monitor links facilitate a communications path to allow the front-panel ports of a single surviving physical chassis segment to remain enabled if the surviving stack segment meets a minimum criteria.

For VSB Stack configurations consisting of more than two 7100-Series devices, LFR monitor links are not used. Instead, LFR uses the cumulative LFR-priority processing described in the sections below.

Whether the front-panel ports of a surviving stack segment remain active or are taken down is based upon the total LFR priority of the surviving stack segment and the total LFR priority of the original VSB stack as follows:

- If the stack segment total LFR priority is greater than half that of the original configured stack total LFR priority, the member switches front-panel ports remain operationally enabled
- If the stack segment total LFR priority is less than half that of the original configured stack, the member switches front-panel ports are disabled

- If the total LFR priority of two stack surviving segments is exactly half of the original configured total, front-panel ports of the stack segment which contains the switch with the highest LFR priority remains operationally enabled

The LFR feature must be globally enabled on each VSB chassis in the VSB system for LFR monitoring to occur. Use the **set bonding lfr enable** command to globally enable LFR on each physical chassis.

The LFR monitor port is configured using the **set bonding port enable** command, the same as a VSB interconnect port. What distinguishes the port types in a VSB context is the port speed. The VSB interconnect port must be a 40GbE port and the LFR monitor port must be either a 1 or 10GbE port.

The LFR priority is configurable on a per chassis basis should you wish to provide a higher likelihood of a chassis or group of chassis belonging to the stack segment for which front-panel ports remain enabled.

A chassis' LFR priority defaults to 10 times the VSB chassis ID. For example, if the VSB chassis ID is 1, the LFR priority is 10. The LFR priority can be manually set using the **set bonding chassis lfr-priority** parameter with a valid range of 1 - 255. Each chassis LFR priority must have a unique value within the stack.

### Example A: LFR Tie Breaker

Example A as presented in [Figure 5-2](#) on page 5-8 depicts a dual failure in a closed interconnect configuration where the two surviving segments results in a LFR priority tie: a priority of 180 for each segment. Chassis 1 – 8 are configured for the default LFR priority of chassis ID \* 10. The total LFR priority for stack segment chassis 1, 2, 7, and 8 is 180. The total LFR priority for stack segment chassis 3 – 6 is 180. The front-panel ports for the stack segment containing chassis 1, 2, 7, and 8 remain enabled because this stack segment contains the chassis with the highest LFR priority: chassis 8, priority 80. The front-panel ports for the stack segment containing chassis 3 – 6 are taken down.

### Example B: Failed Chassis

Example B as presented in [Figure 5-2](#) on page 5-8 depicts both a failure of the bonding link between chassis 1 and 8 as well as a failure of chassis 4. Chassis 1 – 8 are configured for the default LFR priority of chassis ID \* 10. Two stack segments survive. The total LFR priority for the original system is 360. The total LFR priority for stack segment chassis 5 – 8 is 260. Because this stack segment exceeds half the total LFR priority of the original stack, the front-panel ports for the stack segment containing chassis 5 – 8 remain enabled, and the front-panel ports for the stack segment containing chassis 1, 2, and 3 are taken down.

### Example C: Failed Chassis That Takes the Stack Down

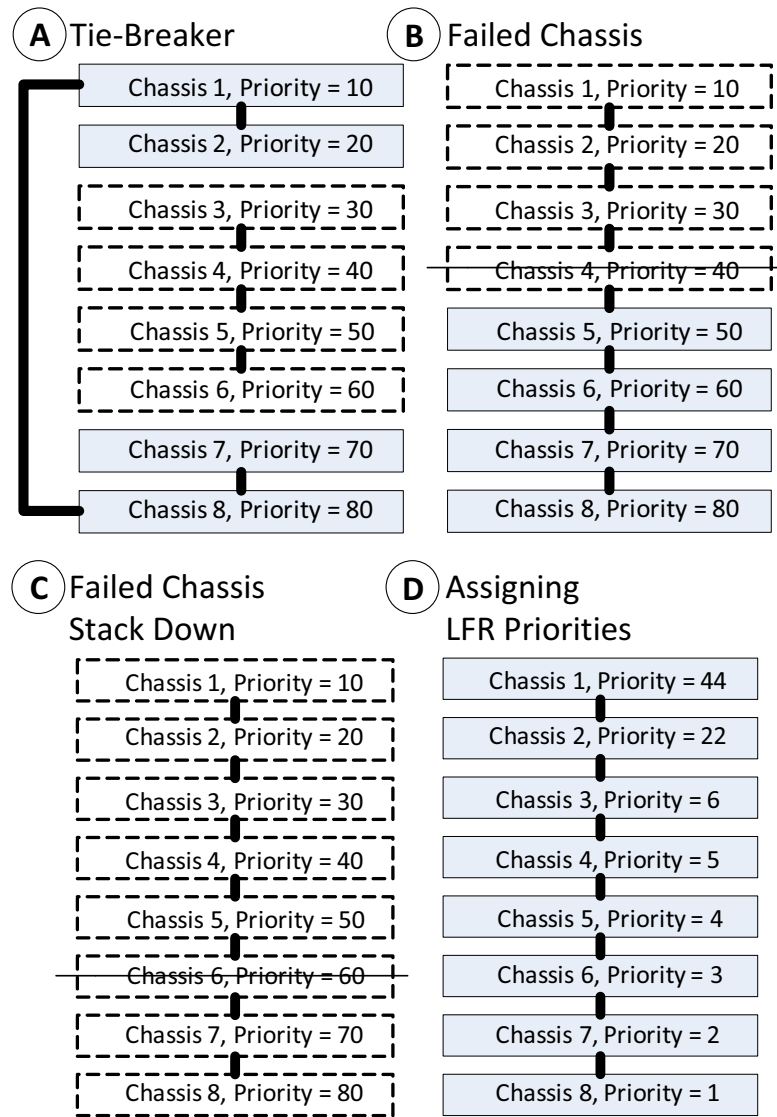
Example C as presented in [Figure 5-2](#) on page 5-8 depicts a chassis failure that takes the entire stack down. In Example B, stack segment 5 – 8 kept its front-panel ports operational because its total LFR priority exceeded half the total LFR priority of the original stack. In this example, each surviving stack LFR priority totals 150. Since the total LFR priority of neither surviving stack segment equaled or exceeded half the total LFR priority value of the original stack (360), the front-panel ports for both stack segments are taken down.

### Example D: Assigning LFR Priorities

This is an example of an LFR priorities administrative configuration. The network administrator wants to assure that if chassis 1 is a member of a surviving stack segment, the front-panel ports of that segment will remain operational regardless of the makeup of the other members. This is true because total LFR priority for the original stack is 87 and the LFR priority for chassis 1 by itself

exceeds half the total LFR priority value the original stack. Because of its LFR priority value, should chassis 1 fail, the front-panel ports for the entire stack is taken down.

**Figure 5-2 VSB LFR Stack Examples**



The VSB feature supports a combined total of 32 VSB 40GbE interconnect and LFR 10GbE monitor links on a VSB system (32 VSB ports per chassis).

Use the **show bonding** command to review LFR state and port configuration details.

## VSB System MAC address Configuration

By default, the VSB system MAC address is set to an internal MAC address associated with VSB chassis 1. You can manually set a MAC address for the VSB system using the **set bonding mac** command. A unique MAC address for each chassis requires a host master election which results in a reset of all non-host master chassis in the stack. This can be disruptive when adding or replacing a chassis in the VSB stack.



**Note:** The VSB system MAC address can not be changed while VSB is globally enabled on the system. You must disable VSB using the `set bonding disable` command before you can modify a globally enabled VSB system MAC address. Disabling VSB on a globally enabled system clears the configuration on both system chassis.

Prior to globally enabling VSB on the system, you can reset a manually configured VSB system MAC address to the default value using the `clear bonding mac` command.

## Licensing

The 7100-Series is auto-entitled for VSB stacking and no license is required to form a VSB stacking system. If you are entitling advanced routing on the VSB stacking system, the advanced routing license must be activated on all chassis in the system. Use the `set license advanced` command, specifying the license key, to configure each chassis with a valid license. The `show license` command displays license status per chassis, before the VSB system is globally enabled, or for all chassis once the VSB system is globally enabled. See “[License Overview](#)” on page 9-10 for license configuration details.

## Globally Enabling and Disabling the VSB System



**Note:** Do not attempt to globally enable the VSB system before:

- Chassis IDs and the VSB system ID are assigned
- At least one interconnect port is VSB enabled between each chassis
- MAC address is assigned, if a non-default MAC address will be used

Chassis configured as bonded chassis with a chassis ID, system ID, and configured with one or more interconnect ports, maintain a status as individual physical chassis until globally enabled for VSB.

Once a chassis is globally enabled for VSB, you can no longer modify VSB configuration for:

- Chassis ID
- System ID
- Mac Address

Fan, PoE and power supply resources are chassis bound. In a VSB globally enabled system, hardware system statistics are displayed per VSB chassis using the `show system hardware` command. PoE capacities are not shared between VSB chassis. PoE resources are assigned to chassis using the `set inlinepower assigned` command and viewed using the `show inlinepower` command, just as you would in a non-VSB system.

Use the `set bonding enable` command to globally enable the VSB system.

Use the `set bonding disable` command to disable a VSB system.



**Note:** Disabling a VSB system clears all non-VSB configuration on both chassis in the system. VSB licensing and a non-enabled VSB configuration persists. If you wish to clear the VSB chassis and system ID configuration after disabling VSB, use the `clear bonding chassis` command.

## High Availability Firmware Upgrade

High Availability firmware Upgrade (HAU) is supported for 7100-Series VSB configurations. Refer to [Chapter 4, High Availability Firmware Upgrade \(HAU\) Configuration](#) for HAU configuration details. Refer to “[High Availability Firmware Upgrade in a Virtual Switch Bonded System](#)” on page 4-6 for HAU configuration details specific to a VSB configuration.

## Applying a VSB Configuration File to a Replacement VSB Chassis

A VSB configuration file contains chassis specific information such as the chassis serial-number. When configuring a replacement VSB chassis with an already existing VSB configuration file, you must specify the chassis ID of the replacement chassis when entering the **configure** command so that chassis specific information on the specified chassis will be ignored and replaced in the configuration file with the correct chassis settings. The chassis ID is specified using the **chassis-id** option.

## Adding a Chassis to an Existing VSB Stacking System

To add an additional chassis in an operating VSB stacking system:

1. Assure that the replacement chassis has been loaded with the same firmware image loaded in the current system.
2. Configure the new chassis' VSB chassis ID to the highest current chassis ID plus 1.
3. Configure the VSB system ID to the system ID of the current stack.
4. Place the new chassis at the bottom of the stack, and cable at least one interconnect link with the chassis above it and one interconnect link with the first chassis in the stack to form a closed ring configuration.
5. Enable bonding on the new chassis' 40GbE interconnect ports.
6. Globally enable the VSB system. The system is reset.
7. When the chassis comes up, VSB will do some processing including rationalizing the MAC address to the system and reboot the device again. The new chassis should then fully join the VSB stacking system.

## A Failed Chassis in an VSB Stacking System

If a chassis fails, there may be traffic disruption lasting several seconds as the system processes link down events and discovers the new topology.

Be aware that it is possible for a stack to be partitioned into separate operational segments. Should this happen, each segment operates as though it is a separate whole stack using the same IP address. The closed ring configuration will prevent separate partitions when a single point of failure occurs by assuring a daisy chain exists for all remaining chassis in the stack. A closed ring configuration is configured using interconnect links between the first and last chassis in the VSB stack.

Topology changes occur when a replacement chassis first joins the stack. Some traffic disruption may accompany these topology changes.

To replace a failed chassis in an VSB stacking system:

1. Remove the failed chassis taking note of the switch/routing, interconnect and uplink cabling that currently exists.
2. Before enabling VSB on the replacement chassis, assure that the replacement chassis has been loaded with the same firmware image loaded in the current system.
3. The configured bonded chassis-ID is associated with the physical chassis in memory. Clear the memory of the failed unit from the stack before configuring bonding for the replacement unit using the **clear bonding chassis *failed-chassis-ID*** command.
4. Configure the replacement chassis' VSB chassis ID and the ID of the VSB system to agree with the configuration of the failed chassis.

5. Optionally, on a two chassis VSB stacking configuration, enable bonding on 10GbE monitor ports.
6. Cable the replacement chassis with the cable configuration noted in step 1.
7. Restore the port level configuration by appending the configuration with the configuration from a previously stored configuration file when the chassis was operational within the stack, using the **configure filename append** command.
8. Enable bonding on the replacement chassis' 40GbE interconnect ports.
9. Globally enable the VSB system. The system is reset.



**Note:** If the bonded MAC address on each chassis is different (as provisioned by the network administrator), a host master election occurs resulting in a reset of all chassis except the elected host master.

## Configuring VSB

This section provides a table of VSB default values and a procedure for configuring a VSB system.

[Table 5-1](#) lists VSB default values.

**Table 5-1 Default VSB Parameters**

Parameter	Description	Default Value
Chassis ID	Identifies a VSB system chassis.	None.
System ID	Identifies a VSB system made up of two VSB chassis.	None.
LFR priority	Used to determine the physical chassis to be put in dormant state should all VSB interconnect links go down.	10 times the chassis ID.
LFR state	Specifies whether LFR is globally enabled or disabled on the physical chassis.	Disabled.
MAC address	VSB system MAC address	An internal MAC address associated with VSB chassis 1.
VSB state	Specifies whether VSB is globally enabled or disabled on the physical chassis.	Disabled.

[Procedure 5-1](#) describes VSB configuration on the Extreme Networks 7100-Series devices. All commands used to configure VSB can be entered in any command mode.

### Procedure 5-1 Configuring VSB

Step	Task	Command(s)
1.	On each chassis, configure each chassis' VSB chassis ID and the ID of the VSB system the two chassis belong to. Optionally: <ul style="list-style-type: none"> <li>• Configure an encrypted secret used on the VSB links</li> <li>• Administratively change the LFR priority for the physical chassis</li> </ul>	<b>set bonding chassis</b> <i>chassis-id</i> { <b>system-id</b> <i>system-id</i>   <b>secret</b> <i>secret</i>   <b>lfr-priority</b> <i>priority</i> }
2.	Optionally change the VSB system MAC address from the default value.	<b>set bonding mac</b> <i>mac-address</i>



**Procedure 5-1 Configuring VSB (continued)**

Step	Task	Command(s)
3.	Enable bonding on the: <ul style="list-style-type: none"> <li>• 40GbE ports to be used as VSB interconnect ports.</li> <li>• 10GbE ports to be used as LFR monitor ports</li> </ul>	<b>set bonding port [port-string] enable</b>
4.	Globally enable LFR on the physical chassis.	<b>set bonding lfr enable</b>
5.	Globally enable the VSB system.	<b>set bonding enable</b>
6.	Optionally configure the VSB system for High Availability firmware upgrades	See “ <a href="#">High Availability Firmware Upgrade</a> ” on page 5-9.

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

## Terms and Definitions

[Table 5-2](#) lists terms and definitions used in this VSB configuration discussion.

**Table 5-2 VSB Configuration Terms and Definitions**

Term	Definition
Virtual Switch Bonding (VSB)	An Extreme Networks 7100-Series feature that aggregates up to eight chassis into a single virtual network device.
Link Failure Response (LFR)	A VSB feature that monitors the VSB interconnect links and determines the physical chassis to be put in a dormant state should all interconnect links go down.
interconnect port	40GbE port that plays the same role as the backplane fabric in a non-VSB chassis by providing distribution between VSB system chassis.
monitor port	10GbE LFR port used by the LFR feature to place the lower priority chassis in a dormant state should all VSB interconnect links go down.
VSB chassis	One of up to eight chassis configured for VSB that make up a VSB system.
VSB system	The aggregation of up to eight chassis configured for VSB and connected by one or more 40GbE interconnect ports.
VSB MAC address	The MAC address for the VSB system that is either manually assigned or defaults to a VSB chassis 1 internal MAC address.

# 6

## Port Configuration

This document describes port configuration on Extreme Networks 7100-Series devices.

For information about...	Refer to page...
<a href="#">Port Configuration Overview</a>	6-1
<a href="#">Configuring Ports</a>	6-12
<a href="#">Terms and Definitions</a>	6-17

### Port Configuration Overview

The Extreme Networks 7100-Series modules and standalone devices have fixed front panel switch ports and, depending on the model, optional expansion module slots. The numbering scheme used to identify the switch ports on the front panel and the expansion module(s) installed is interface-type dependent and is also dependent upon the chassis slots in which the module(s) are installed. Port numbering proceeds from 1 to the maximum number of that port type on the module. If there are multiple port types, each port type numbering starts at 1. Port numbering is displayed next to each port.

When configuring a port, the port string associated with a port is made up of the port type, the slot location of the module in the chassis and the port number delineated by a period as explained in [Port String Syntax Used in the CLI](#).

The following topics are covered in this section:

For information about...	Refer to page...
<a href="#">Port String Syntax Used in the CLI</a>	6-2
<a href="#">Console Port Parameters</a>	6-3
<a href="#">Administratively Enabling a Port</a>	6-4
<a href="#">Ingress Filtering</a>	6-4
<a href="#">Port Alias</a>	6-4
<a href="#">Force Linkdown</a>	6-4
<a href="#">Default Port Speed</a>	6-5
<a href="#">Port Duplex</a>	6-5
<a href="#">Jumbo Frames</a>	6-6
<a href="#">Auto-Negotiation and Port Advertised Ability</a>	6-7
<a href="#">Port MDI/MDIX</a>	6-8

For information about...	Refer to page...
<a href="#">Port Flow Control</a>	6-8
<a href="#">Configuring Link Traps and Link Flap Detection</a>	6-9
<a href="#">Port Broadcast Suppression</a>	6-10
<a href="#">Port Priority</a>	6-10
<a href="#">Transmit Queue Monitoring</a>	6-11

## Port String Syntax Used in the CLI

Commands requiring a *port-string* parameter use the following syntax to designate the type of port being configured, slot location the module containing the port is inserted into the chassis, and port number on the module containing the port:

**port type.slot location.port number**

Where **port type** can be:

- **ge** - 1-Gbps Ethernet
- **tg** - 10-Gbps Ethernet
- **fg** - 40-Gbps Ethernet
- **com** - COM (console) port
- **host** - the host port
- **vlan** - VLAN interfaces
- **lag** - IEEE802.3 link aggregation ports
- **lpbk** - loopback interfaces, or
- **lo** - the local (software loopback) interface
- **vtap** - a MIB-II interface for VLANs, used as the data source input of a port mirror or SMON statistics collection on that particular VLAN

**Slot location for modules installed in a 7100-Series chassis can be:**

0 through the maximum number of slots in the chassis, with 0 designating virtual system ports (lag, vlan, host, loopback), and 1 designating the lowest module slot in the chassis.

**Port number** can be:

Any port number on the module. The highest valid port number is dependent on the number of ports in a slot location and the port type.

**For example:**

If the 7100-Series device has 48 10GbE front panel ports and 4 40GbE uplink ports, the range of port number designations used in the CLI command would be:

tg.1.1 through tg.1.48 for the 48 10GbE front panel ports, and  
fg.1.1 through fg.1.4 for the 4 40GbE uplink ports.

### Examples



**Note:** You can use a wildcard (\*) to indicate all of an item. For example, tg.1.\* would represent all 10GbE ports in the module in slot 1.

This example shows the *port-string* syntax for specifying ports 1, 3 and 11 in the module in chassis slot 1:

```
tg.1.1;tg.1.3;tg.1.11
```

This example shows the *port-string* syntax for specifying ports 1, 3, 7, 8, 9 and 10 in the module in chassis slot 1:

```
tg.1.1,tg.1.3,tg.1.7-10
```

This example shows the *port-string* syntax for specifying the 40-GbE port 2 of the module in chassis slot 1:

```
tg.1.2
```

This example shows the *port-string* syntax for specifying all 10GbE ports in the module in chassis slot 1:

```
tg.1.*
```

This example shows the *port-string* syntax for specifying all 10-Gbps Ethernet ports in the chassis:

```
tg.*.*
```

This example shows the *port-string* syntax for specifying all 40-Gbps Ethernet ports in the chassis:

```
fg.*.*
```

This example shows the *port-string* syntax for specifying all ports (of any interface type) in all modules in the chassis:

```
*.*.*
```

## Console Port Parameters

Each Extreme Networks 7100-Series module or standalone device includes a console port through which local management of the device can be accessed using a terminal or modem. The CLI provides for:

- The display of console port configurations using the **show console** command in any command mode
- The setting of console port parameters, including the baud rate, flow control, number of bits, number of stop bits and parity, using the **set console** command in any command mode
- The clearing of console port parameters to default values using the **clear console** command in any command mode

If C2 security mode is enabled, You can not create, modify, or clear a console configuration while in Read-Write user mode.

When specifying a console port string, use the **com** keyword for the port type, as specified in the [Port String Syntax Used in the CLI](#) discussion.

The following example shows how to set the baud rate to 19200 on console port com.1.1:

```
7100 Chassis(rw)->set console baud 19200 com.1.1
```

The following example shows how to set the bits property value to 8 on all console ports:

```
7100 Chassis(rw)->set console bits 8
```

The following example shows how to set the flowcontrol property value to none on console port com.1.1:

```
7100 Chassis(rw)->set console flowcontrol none com.1.1
```

The following example shows how to set the parity property value to even on all ports:

```
7100 Chassis(rw)->set console parity even
```

The following example shows how to set the stopbits property value to one on console ports com.1.1 and com.1.2:

```
7100 Chassis(rw)->set console stopbits one com.1.1-2
```

## Administratively Enabling a Port

Ports are administratively disabled by default.

Use the **set port enable** command to administratively enable the specified ports.

Use the **set port disable** command to administratively disable the specified ports.

The following example administratively enables port tg.1.1:

```
7100 Chassis(rw)->set port enable tg.1.1
```

```
7100 Chassis(rw)->show port tg.1.1
```

```
Port tg.1.1 enabled
```

```
7100 Chassis(rw)->
```

## Ingress Filtering

The ingress filtering feature provides for a means of limiting the forwarding of received frames on the ingress port based on the VLAN egress list for that port. VLAN IDs of a port's incoming frames are compared to the port's egress list. If the received VLAN ID does not match a VLAN ID on the port's egress list, the frame is dropped.

Use the **set port ingress-filter** command in any command mode to enable ingress filtering on the specified ports.

The following example enables ingress filtering on port tg.1.1

```
7100 Chassis(rw)->set port ingress-filter tg.1.1 enable
```

```
7100 Chassis(rw)->>show port ingress-filter tg.1.1
```

```
Port          State
-----
tg.1.1        enabled
```

## Port Alias

The alias feature allows a string name to be associated with a port.

Use the **set port alias** command to configure an alias for the specified ports.

The following example sets the alias on port tg.1.1 to **documentation**

```
7100 Chassis(rw)->set port alias tg.1.1 documentation
```

```
7100 Chassis(rw)->>show port alias tg.1.1
```

```
Alias on port tg.1.1 set to: Documentation.
```

## Force Linkdown

When the force linkdown feature is disabled, disabling a port using **set port disable** will disable the ability to forward traffic, but the link stays up. When force linkdown is enabled, disabling a port using **set port disable** will disable the link completely.

When force linkdown is enabled, disabling a port using the **set port disable** command will not disable PoE on that port.

Force linkdown is disabled by default.

Use the **set forcmlinkdown** command in any command mode to enable the force linkdown feature on this device.

The following example enables the force linkdown feature on this device:

```
7100 Chassis(rw)->set forcmlinkdown enable
7100 Chassis(rw)->show forcmlinkdown
ForceLinkDown feature is globally enabled.
7100 Chassis(rw)->
```

## Default Port Speed

When auto-negotiation is enabled, the port speed used is determined by the fastest compatible speed between linked ports. On ports capable of multiple speeds, if auto-negotiation is not enabled, the default port speed setting provides for the configuring of a default speed for this port. Use the **set port speed** command to specify the default speed for the specified ports. Valid values are 1000, 10000, and 40000 Mbps.

When a 40GB QSFP fiber MAU is present, the system cannot determine from the QSFP cabling whether the mode should be set to 1x40GB mode or 4x10GB mode. The port defaults to 1x40GB mode. To set 4x10GB mode, enter **set port speed port-string 10000**. Use the 40GB **fg.x.y** designation for the port-string. The port will move to a down state and remain down until the system is reset. After the system reset, the port will come back up as four 10GB ports using the port-string format **tg.x.y**.

To set the 4x10GB mode back to 1x40GB mode, enter the command **set port speed port-string 40000** specifying one of the 10GB ports. The four 10GB ports will move to a down state and remain down until the system is reset. After the system reset, the port will come back up as a single 40GB port using the port-string format **fg.x.y**.

Auto-negotiation is enabled by default.

The following example sets the default speed on port tg.1.1 to 1000 Mbps:

```
7100 Chassis(rw)->set port speed tg.1.1 1000
7100 Chassis(rw)->show port speed tg.1.1
default speed is 1000 on port tg.1.1.
7100 Chassis(rw)->
```

## Port Duplex

Duplex between two communicating devices specifies whether communication will be one way at a time (half-duplex) or in both directions simultaneously (full-duplex). When auto-negotiation is enabled, auto-negotiation determines port duplex.

Use the **set port duplex** command to specify whether the specified ports will operate at half or full duplex when auto-negotiation is not enabled.

The following example sets the port duplex on port tg.1.1 to full:

```
7100 Chassis(rw)->set port duplex tg.1.1 full
7100 Chassis(rw)->show port duplex tg.1.1
default duplex mode is full on port tg.1.1.
```

7100 Chassis (rw) ->

## Jumbo Frames

The jumbo frames feature supports Ethernet frames greater than 1500 bytes of payload on a port. By default, jumbo frame support is disabled on all ports and path MTU discovery is enabled. When jumbo frame support is enabled, path MTU discovery should also be enabled. Path MTU discovery is set using the **set mtu** command.

It is possible for jumbo administrative status to be enabled and jumbo operational status to be deferred. Jumbo frame support is supported on all module ports.

When enabling jumbo frame support, the maximum frame size defaults to 10239 bytes for untagged packets and supports 10243 bytes for tagged packets. Jumbo frame size can be set to any value between 1000 – 10239 bytes. The standard Ethernet frame MTU for untagged packets is 1518 bytes including the 18 Ethernet header bytes. The standard Ethernet frame MTU for tagged packets is 1522 bytes.

When setting the jumbo frame size, keep in mind that large frames require more packet buffers than standard frames, reducing the total available packet buffers. This limitation can be improved by selecting a smaller MTU jumbo frame size.

Some applications require extra header bytes beyond the standard tagged packet size. For example, provider bridging (Q-in-Q) requires an extra header of 4 additional bytes. The 7100-Series supports the dynamic MTU frame which when enabled will automatically add the extra 4 bytes to the currently configured jumbo MTU frame size, allowing you to set the jumbo MTU frame size to the standard 1518 bytes and the 4 additional tagged header bytes will automatically be added for Q-in-Q frames. The dynamic MTU frame feature currently only supports Q-in-Q. Other applications requiring additional header bytes will be added in future releases. Dynamic MTU frame is disabled by default.

Use the **show port jumbo** command to verify the operational status of a jumbo enabled port including display of supported jumbo MTU frame size and the state of the dynamic MTU frame feature. A jumbo administratively disabled port will always have a jumbo operational status of disabled.

The 7100-Series platform supports the forwarding of unicast, multicast, and broadcast jumbo frames.

Use the **set port jumbo** command in any command mode to enable or disable jumbo frame support on the specified ports.

Use the **set port jumbo mtu** command in any command mode to change the size of the jumbo MTU frame on the port. Jumbo MTU frame must be enabled for the frame size change to take affect.

Use the **set port jumbo mtu dynamic** command to enable or disable dynamic MTU frame on the port.

Use the **show port jumbo** command to verify the operational status of a jumbo enabled port.

The following example enables the port jumbo frame feature, sets the jumbo MTU frame size to 1522 and enables dynamic MTU frame on port tg.1.1 (Q-in-Q must be enabled for the MTU delta and Application name fields to display as presented here):

```
7100 Chassis(rw)->set port jumbo enable tg.1.1
7100 Chassis(rw)->set port jumbo mtu 1522 tg.1.1
7100 Chassis(rw)->set port jumbo mtu dynamic enable tg.1.1
7100 Chassis(rw)->show port jumbo tg.1.20
```

\* Applicable only if port jumbo is enabled

Port	Oper Status	Admin Status	MTU Dynamic	MTU Min	MTU Max	MTU Oper*	MTU Admin	MTU Delta	Application Name (size)
tg.1.1	Enabled	Enabled	Enabled	1000	10239	1522	1522	4	Q-in-Q

7100 Chassis(rw) ->

## Auto-Negotiation and Port Advertised Ability

Auto-negotiation is an Ethernet feature that facilitates the selection of port speed, duplex, and flow control between the two members of a link, by first sharing these capabilities and then selecting the fastest transmission mode that both ends of the link support. Auto-negotiation is enabled by default.

The advertised ability feature allows for the port to share its port capabilities with the other end of the link. Advertised capabilities will be used during the auto-negotiation process. Actual port capabilities, advertised port capability and remote end advertised port capabilities can be displayed using the **show port advertise** command in any command mode. The following port capabilities can be advertised:

- **10t** - 10BASE-T half duplex mode
- **10tfd** - 10BASE-T full duplex mode
- **100tx** - 100BASE-TX half duplex mode
- **100txfd** - 100BASE-TX full duplex mode
- **1000x** - 1000BASE-X, -LX, -SX, -CX half duplex mode
- **1000xfd** - 1000BASE-X, -LX, -SX, -CX full duplex mode
- **1000t** - 1000BASE-T half duplex mode
- **1000tfd** - 1000BASE-T full duplex mode
- **pause** - PAUSE for full-duplex links
- **apause** - Asymmetric PAUSE for full-duplex links
- **spause** - Symmetric PAUSE for full-duplex links
- **bpause** - Asymmetric and Symmetric PAUSE for full-duplex links



**Note:** Advertised ability can be activated only on ports that have auto-negotiation enabled.

During auto-negotiation, making use of information gained from the advertised ability feature, the port “tells” the device at the other end of the segment what its capabilities and mode of operation are. If auto-negotiation is disabled, the port reverts to the values specified by the default speed, default duplex, and the port flow control commands.

Use the **set port negotiation** command to enable auto-negotiation on the specified ports.

Use the **set port advertise** command to specify the capabilities to be advertised on the specified ports.

The following example enables auto-negotiation on port tg.1.1 and sets the advertise utility to advertise 10BASE-T half duplex mode, 10BASE-T full duplex mode, 100BASE-TX half duplex



mode, 100BASE-TX full duplex mode, and Asymmetric and Symmetric PAUSE for full-duplex links:

```
7100 Chassis(rw)->set port negotiation tg.1.1 enable
7100 Chassis(rw)->set port advertise tg.1.1 10t 10tfd 100tx 100txfd bpause
7100 Chassis(rw)->show port advertise tg.1.1
tg.1.1      capability  advertised  remote
-----
10BASE-T    yes           yes        yes
10BASE-TFD  yes           yes        yes
100BASE-TX  yes           yes        yes
100BASE-TXFD yes          yes        yes
1000BASE-X  no            no         no
1000BASE-XFD no            no         no
1000BASE-T  no            no         no
1000BASE-TFD no            no         no
other       no            no         no
pause       yes           no         yes
Apause      yes           no         no
Spause      yes           no         yes
Bpause      yes           yes        no
7100 Chassis(rw)->
```

## Port MDI/MDIX

The Port MDI/MDIX feature detects and adapts to straight through (MDI) or cross-over (MDIX) Ethernet cabling on switch ports. Ports can be set to auto detect, force MDI or force MDIX. The default is for auto-detection of the cabling type.

Use the **set port mdix** command in any command mode to set the MDI/MDIX feature for the specified ports on this device.

The following example sets the MDI/MDIX feature to cross-over for all ports on this device.

```
7100 Chassis(rw)->set port mdix mdix
7100 Chassis(rw)->
```

## Port Flow Control

Flow control is the process of managing the rate of data transmission between two nodes to prevent a fast sender from overrunning a slow receiver. It provides a mechanism for the receiver to control the transmission speed. Flow control helps prevent congestion. Flow control should be distinguished from congestion control, which is used for controlling the flow of data when congestion has actually occurred. Flow control on a port is configured for whether the port sends, receives or both sends and receives flow control packets.

When auto-negotiation is enabled the port flow control settings have no bearing on flow control. Pause is negotiated through the predefined advertised settings. The port flow control settings take effect when auto-negotiation is disabled.

Use the **set port flowcontrol** command to both enable flow control and configure the flow control setting for all or the specified ports.

The following example sets flow control on port tg.1.1 to both send and receive flow control packets:

```
7100 Chassis(rw)->set port flowcontrol tg.1.1 both enable
7100 Chassis(rw)->
```

## Configuring Link Traps and Link Flap Detection

The link traps and link flap detection features provide for the disabling or re-enabling of link traps and to configure the link flapping detection function. By default, all ports are enabled to send SNMP trap messages indicating changes in their link status (up or down).

Use the `set port trap` command in any command mode to enable the sending of SNMP trap messages when link status changes.

The following example enables SNMP traps on port tg.1.1:

```
7100 Chassis(rw)->set port trap tg.1.1 enable
7100 Chassis(rw)->show port trap tg.1.1
Link traps enabled on port tg.1.1.
7100 Chassis(rw)->
```

The link flap function detects when a link is going up and down rapidly (also called “link flapping”) on a physical port, and takes the configured actions (disable port, and eventually send notification trap) to stop such a condition. If left unresolved, the “link flapping” condition can be detrimental to network stability because it can trigger Spanning Tree and routing table recalculation.

The link flap utility is disabled both globally and on ports by default. The link flap utility must be enabled globally and on the ports for which link flap detection is to occur.

Use the **set linkflap globalstate** command in any command mode to globally enable the link flap utility on this device.

Use the **set linkflap portstate** command in any command mode to enable the link flap utility on the specified ports.

There are three link flap actions that can be configured as a response to link flapping:

- Disable the interface
- Generate a SYSLOG message
- Generate an SNMP trap

You can also set the action to all three. A link flap action will occur if the number of link flaps exceeds the configured link flap threshold (number of times the link flaps) setting within the period configured by link flap interval.

Use the **set linkflap action** command in any command mode to set the link flap action for the specified ports.

Use the **set linkflap threshold** command in any command mode to set the number of link flaps that will trigger a link flap action for the specified ports.

Use the **set linkflap interval** command in any command mode to set the period of time within which the link flap threshold must be exceeded to cause the link flap action to trigger.

If the link flap action is to disable the interface, a port downtime period in seconds can be configured to specify how long the disabled interface will remain down. A value of 0 indicates forever.

Use the **set linkflap downtime** command in any command mode to configure the downtime period for the specified ports.

The following example configures the link flap utility on port `tg.1.1` to:

- Set the link flap action to all three actions
- Set the link flap threshold to 12 link flaps
- Sets the link flap interval to 6 seconds
- Sets the downtime period to 600 seconds

```
7100 Chassis(rw)->set linkflap action tg.1.1 all
7100 Chassis(rw)->set linkflap threshold tg.1.1 12
7100 Chassis(rw)->set linkflap interval tg.1.1 6
7100 Chassis(rw)->set linkflap downtime tg.1.1 600
7100 Chassis(rw)->show linkflap parameters tg.1.1
Linkflap Port Settable Parameter Table (X means error occurred)
Port      LF Status  Actions  Threshold  Interval  Downtime
-----  -
tg.1.1    disabled  D..S..T  12         6         600
1 port(s) found.
7100 Chassis(rw)->
```

## Port Broadcast Suppression

Broadcast suppression sets a threshold on the broadcast traffic that is received and switched out to other ports. The maximum value in packets per second is 1488100. If the maximum value is configured, broadcast suppression is disabled. Broadcast suppression is disabled by default.

Use the **set port broadcast** command in any command mode to set the broadcast suppression limit, in packets per second, on the specified ports.

The following example sets the broadcast suppression threshold to 10000 packets per second for port `tg.1.1`:

```
7100 Chassis(rw)->set port broadcast tg.1.1 10000
7100 Chassis(rw)->show port broadcast tg.1.1
Port      Total BC      Threshold      Peak Rate      Peak Rate Time
          Packets      (pkts/s)      (pkts/s)      (ddd:hh:mm:ss)
-----  -
tg.1.1    784628        10000         2400          000:00:02:11
```

## Port Priority

The Extreme Networks 7100-Series device supports Class of Service (CoS), which allows you to assign mission-critical data to higher priority through the device by delaying less critical traffic during periods of congestion. The higher priority traffic through the device is serviced first before lower priority traffic. The Class of Service capability of the device is implemented by a priority queueing mechanism. Class of Service is based on the IEEE 802.1D (802.1p) standard specification, and allows you to define eight priorities (0 through 7) and, depending on port type, up to 16 transmit queues (0-15) of traffic for each port.

A priority 0 through 7 can be set on each port, with 0 being the lowest priority. A port receiving a frame without priority information in its tag header is assigned a priority according to the default

priority setting on the port. For example, if the priority of a port is set to 4, the frames received through that port without a priority indicated in their tag header are classified as a priority 4 and transmitted according to that priority.

In addition, the device's rate limiting capabilities allow you to further prioritize traffic by limiting the rate of inbound or outbound traffic on a per port/priority basis.



**Note:** When CoS override is enabled using the **set policy profile** command as described in the “Policy Profile Commands” section of the *Extreme Networks 7100-Series CLI Reference*, CoS-based classification rules will take precedence over priority settings configured with the **set port priority** command described in this section.

Use the **set port priority** command in any command mode to set the port priority for the specified ports.

The following example sets the port priority for port tg.1.1 to 4:

```
7100 Chassis(rw)->set port priority tg.1.1 4
7100 Chassis(rw)->show port priority tg.1.1
tg.1.1 is set to 4
```

## Transmit Queue Monitoring

The transmit queue monitoring feature monitors ports for stalled buffers. 7100-Series firmware releases 7.91.03 and greater support transmit queue monitoring. When a port is congested, if a configured minimum number of packets are not transmitted within the sample interval set in seconds, the transmit buffer is considered to be in a stalled state. The stalled buffer condition can occur for reasons such as duplex misconfiguration or the continuous reception of pause frames that prevent the port from transmitting. If no intervention takes place, it is possible for all packet buffers to stall preventing any new traffic from being received or forwarded.

When transmit queue monitoring detects a stalled buffer, a set of three independently configured actions will occur if a configured error threshold of sample intervals for each action is met:

- The logging threshold generates a Syslog message notification that the condition exists
- The ignore pause threshold disables processing of received pause packets on the port
- The shutdown port threshold disables the port



**Note:** Setting the disable port threshold to trigger before the ignore pause frames or logging thresholds causes these action thresholds to never be met and the action will not occur.

The ignore pause and disable port thresholds support a configured hold time. The hold time is the amount of time the port will remain in the ignore pause or disabled port state. Once the hold time expires, the port returns to the normal state and continues to monitor the stalled condition. If the hold time is set to 0, the port will not return to the normal state until you administratively clear the condition.

The transmit queue monitoring feature is enabled by default and can be disabled or re-enabled using the **set txqmonitor state** command.

The sample interval, during which a configured minimum number of packets must be transmitted by the port buffer to avoid being set to the buffer stalled state, can be configured in seconds using the **set txqmonitor sampleinterval** command. The sample interval defaults to 1 second.

The minimum number of packets that must be transmitted by the transmit buffer to avoid being set to the buffer stalled state can be configured using the **set txqmonitor minrate** command. The minimum number of packets transmitted by the transmit buffer defaults to 1 packet.

The action thresholds in number of consecutive errored sample intervals for logging, ignoring pause frames, and disabling the port can be configured using the **set txqmonitor threshold** command. The logging threshold is 2 consecutive errored sample intervals. The ignore pause frames threshold is 5 consecutive errored sample intervals. The disable port threshold is 10 consecutive errored sample intervals.

The amount of time a port will be down if the disable port threshold action is triggered can be set using the **set txqmonitor downtime** command. The default period is set to 0 sample intervals. The port must be administratively cleared using the **clear txqmonitor operstatus** command.

The amount of time a port will ignore all pause frames it receives if the ignore pause action threshold is met can be configured using the **set txqmonitor ignorepausetime** command. The default period is set to 0 sample intervals. The ignore pause time must be administratively cleared using the **clear txqmonitor operstatus** command.

The generation of traps each time a transmit queue monitoring threshold is met can be enabled or disabled using the **set txqmonitor trapstatus** command. The generation of traps is enabled by default.

## Energy Efficient Ethernet (EEE)

Enabling EEE on a link reduces the power consumption on the Ethernet link during low data activity. EEE must be enabled on both sides of the link to operate. Auto negotiation is restarted when EEE is enabled or disabled, causing the link to bounce. Link state does not change as a result of an EEE transition to and from a lower level of power. Frames that are in transit are neither dropped nor corrupted during EEE transition to and from a lower level of power.

The wakeup time is the period between the reception of an IDLE signal and the reception of the first data permitted on the interface. It is recommended that you only modify wakeup and fallback values if a longer wakeup time is required. The negotiation of wakeup times is accomplished using the LLDP EEE TLV which must be enabled on both sides of the link using the **set lldp port tx-tlv energy-eff-eth** command. If the configured wakeup time is not acceptable, a fallback wakeup time is used.

Use the **set port energy-eff-eth** command to enable EEE and change wakeup and fallback times.

## Configuring Ports

This section provides details for the configuration of ports on the 7100-Series products.

[Table 6-1](#) lists port parameters and their default values.

**Table 6-1 Default Port Parameters**

Parameter	Description	Default Value
broadcast suppression	Specifies a limit for the number of broadcast packets per second that can be received and switched on a port.	disabled (set to max value of 1488100)
console baud rate	Specifies the baud rate for the console port.	9600
console bits	Specifies the number of bits per character on the console port.	8 bits
console flow control	Specifies the flow control mechanism for the console port.	ctsrts (Clear to Send/Request to Send)

**Table 6-1 Default Port Parameters (continued)**

Parameter	Description	Default Value
global link flap state	Specifies whether link flap is enabled globally on this device	disabled
jumbo frame support	Specifies whether Ethernet frame with a payload greater than 1500 is supported on this port.	disabled
port ingress filter	Specifies that frame forwarding is limited to members of the port's VLAN egress list.	disabled
port negotiation	Specifies whether auto-negotiation is enabled on this port.	enabled
port priority	Specifies the 802.1D priority for this port.	0
port state	Specifies the port state.	disabled
port traps	Specifies whether the sending of port traps is enabled on this port.	enabled
transmit queue monitoring state	Specifies the global state of transmit queue monitoring on the device.	enabled
transmit queue monitoring trap status	Specifies the global state of transmit queue monitoring trap generation on the device.	enabled
transmit queue monitoring down time	Specifies the number of sample intervals the port will be down if the disable port threshold action is triggered for the port.	1 sample interval
transmit queue monitoring ignore pause frame time	Specifies the number of sample intervals the port will ignore pause frames if the the ignore pause frames threshold action is triggered on the port.	1 sample interval
transmit queue monitoring minimum packet sample rate	Specifies the number of packets transmitted by the transmit buffer within the sample interval to avoid the buffer being set to the stalled state.	1 packet.
transmit queue monitoring logging action threshold	Specifies the logging action threshold in sample intervals before transmit queue monitoring logs an error.	2 sample intervals
transmit queue monitoring ignore pause frames action threshold	Specifies the ignore pause frames action threshold in sample intervals before the ignore pause frame action is triggered.	5 sample intervals
transmit queue monitoring disable port action threshold	Specifies the disable port action threshold in sample intervals before the port disable action is triggered.	10 sample intervals
port buffer mode	Specifies the port buffer mode setting for the device.	flow-control

[Procedure 6-1](#) describes how to configure ports.

### Procedure 6-1 Configuring Ports

Step	Task	Command(s)
1.	Administratively enable one or more ports on the system.	<b>set port enable</b> <i>port-string</i>
2.	Optionally, change the properties for one or more console ports.	<b>set console</b> {[ <i>baud rate</i> ]   [ <i>bits num-bits</i> ]   [ <i>cts-link</i> { <i>enable</i>   <i>disable</i> }]   [ <i>flowcontrol</i> { <i>none</i>   <i>ctsrts</i>   <i>dsrdtr</i> }]   [ <i>parity</i> { <i>none</i>   <i>odd</i>   <i>even</i>   <i>mark</i>   <i>space</i> }]   [ <i>stopbits</i> { <i>one</i>   <i>oneandhalf</i>   <i>two</i> }]   [ <i>vt100 dsr</i> { <i>enable</i>   <i>disable</i>   <i>timeout timeout</i> }] [ <i>port-string</i> ]
3.	Optionally, limit the forwarding of received frames based on port VLAN egress lists.	<b>set port ingress-filter</b> <i>port-string</i> <b>enable</b>
4.	Optionally, assign an alias name to a port.	<b>set port alias</b> <i>port-string</i> [ <i>string</i> ]
5.	Optionally, change the port buffer mode.	<b>set port buffer mode</b> { <i>flow-control</i>   <i>priority-groups</i> }
6.	Optionally, enable the forcing of ports in the “operstatus down” state to become disabled.	<b>set forcelinkdown</b> <b>enable</b>
7.	Optionally, set the default speed of one or more ports.	<b>set port speed</b> <i>port-string</i> { <b>10</b>   <b>100</b>   <b>1000</b>   <b>10000</b>   <b>40000</b> }
8.	Optionally, set the default duplex type for one or more ports.	<b>set port duplex</b> <i>port-string</i> { <b>full</b>   <b>half</b> }
9.	Optionally, enable jumbo frame support on one or more ports.	<b>set port jumbo</b> <b>enable</b> [ <i>port-string</i> ]
10.	Optionally, enable auto-negotiation on one or more ports.	<b>set port negotiation</b> <i>port-string</i> <b>enable</b>
11.	Optionally, set MDI/MDIX mode on one or more ports.	<b>set port mdix</b> [ <i>port-string</i> ] { <b>auto</b>   <b>mdi</b>   <b>mdix</b> }
12.	Optionally, configure the auto-negotiation advertised capabilities on one or more ports.	<b>set port advertise</b> <i>port-string</i> {[ <b>10t</b> ] [ <b>10tfd</b> ] [ <b>100tx</b> ] [ <b>100txfd</b> ] [ <b>1000x</b> ] [ <b>1000xfd</b> ] [ <b>1000t</b> ] [ <b>1000tfd</b> ] [ <i>pause</i> ] [ <i>apause</i> ] [ <i>spause</i> ] [ <i>bpause</i> ]}
13.	Optionally, enable flow control settings for one or more ports.	<b>set port flowcontrol</b> <i>port-string</i> { <b>receive</b>   <b>send</b>   <b>both</b> } <b>enable</b>
14.	Optionally, set the broadcast suppression limit on one or more ports.	<b>set port broadcast</b> <i>port-string</i> <i>threshold-val</i>
15.	Optionally, set a default port priority for one or more ports.	<b>set port priority</b> <i>port-string</i> <i>priority</i>

[Procedure 6-2](#) describes how to configure transmit queue monitoring.

### Procedure 6-2 Configuring Transmit Queue Monitoring

Step	Task	Command(s)
1.	Enable transmit queue monitoring globally on the device if it is disabled.	<b>set txqmonitor state</b> { <b>enable</b>   <b>disable</b> }
2.	Optionally, modify the transmit queue monitoring sample interval.	<b>set txqmonitor sampleinterval</b> <i>seconds</i>

**Procedure 6-2 Configuring Transmit Queue Monitoring (continued)**

Step	Task	Command(s)
3.	Optionally, modify the number of sample intervals the port will be down if the disable port threshold action is triggered.	<b>set txqmonitor downtime</b> <i>sample-intervals</i>
4.	Optionally, modify the number of sample-intervals the port will ignore pause frames if the ignore pause frames threshold action is triggered.	<b>set txqmonitor ignorepausetime</b> <i>sample-intervals</i>
5.	Optionally, modify the minimum number of packets transmitted by the transmit buffer within the sample interval to avoid the buffer being set to stalled state.	<b>set txqmonitor minrate</b> <i>packets</i>
6.	Optionally, modify the number of sample intervals before a transmit queue monitoring action is triggered.	<b>set txqmonitor threshold</b> { <b>disableport</b>   <b>ignorepause</b>   <b>logging</b> } <i>sample-intervals</i>
7.	Optionally, enable transmit queue monitoring global traps generation.	<b>set txqmonitor trapstatus</b> { <b>enable</b>   <b>disable</b> }

[Procedure 6-3](#) describes how to configure link trap and link flap detection.

**Procedure 6-3 Configuring Link Trap and Link Flap Detection**

Step	Task	Command(s)
1.	Optionally, enable one or more ports for sending SNMP trap messages when link status changes occur.	<b>set port trap</b> <i>port-string</i> <b>enable</b>
2.	Optionally, globally enable the link flap detection function for this device. Defaults to disabled.	<b>set linkflap globalstate</b> <b>enable</b>
3.	Optionally, enable the link flap detection function on one or more ports. Defaults to disabled.	<b>set linkflap portstate</b> <b>enable</b> [ <i>port-string</i> ]
4.	Optionally, change the period of time within which the link flap threshold must be exceeded to cause the link flap action to trigger.	<b>set linkflap interval</b> <i>port-string</i> <i>interval_value</i>
5.	Optionally, set the action that will occur when a link flap violation threshold is met.	<b>set linkflap action</b> <i>port-string</i> { <b>disableinterface</b>   <b>gensyslogentry</b>   <b>gentrap</b>   <b>all</b> }
6.	Optionally, change the link flap action trigger threshold.	<b>set linkflap threshold</b> <i>port-string</i> <i>threshold_value</i>
7.	Optionally, set the length of time one or more ports will be held down after a link flap violation threshold is met and the action is set to disable the interface.	<b>set linkflap downtime</b> <i>port-string</i> <i>downtime_value</i>



Table 6-2 describes how to manage port configuration.

**Table 6-2 Managing Port Configuration**

Task	Command
To clear the properties set for one or more console ports to its default values:	<b>clear console</b> [baud] [bits] [cts-link] [flowcontrol] [parity] [stopbits] [vt100] [port-string]
To override the causes configured to place operating status to a down or dormant state for one or more ports:	<b>clear port operstatuscause</b> [port-string] [admin] [all] [cos] [flowlimit] [linkflap] [policy]
To reset the force link down function to the default state of disabled:	<b>clear forcelinkdown</b>
To reset jumbo frame support status to enabled on one or more ports:	<b>clear port jumbo</b> [port-string]
To reset MDIX mode to the default setting of auto on one or more ports:	<b>clear port mdix</b> [port-string]
To reset auto-negotiation advertised capabilities to the default setting on one or more ports:	<b>clear port advertise</b> port-string [10t   10tfd   100tx   100txfd   1000x   1000txfd   1000t   1000tfd   pause   apause   spause   bpause]
To clear the configured actions to a link flap violation:	<b>clear linkflap action</b> {port-string} {disableInterface   gensyslogentry   gentrap   all}
To toggle link flap disabled ports to operational:	<b>clear linkflap down</b> [port-string]
To clear all link flap options or statistics on one or more ports:	<b>clear linkflap</b> {all   stats [port-string]   parameter port-string {threshold   interval   downtime   all}}
To reset the broadcast threshold or clear the peak rate and peak time values on one or more ports:	<b>clear port broadcast</b> port-string {[threshold] [peak]}
To reset the current default port priority setting to the default value of 0 on one or more ports:	<b>clear port priority</b> port-string

Table 6-3 describes how to display port configuration information and statistics.

**Table 6-3 Displaying Port Configuration Information and Statistics**

Task	Command
To display properties set for one or more console ports:	<b>show console</b> [baud] [bits] [flowcontrol] [parity] [stopbits] [port-string]
To display whether or not one or more ports are enabled for switching:	<b>show port</b> [port-string]
To display operating and admin status, speed, duplex mode and port type for one or more ports on the device:	<b>show port status</b> [port-string] [-interesting]
To display port counter statistics detailing traffic through the device and through all MIB2 network devices:	<b>show port counters</b> [port-string] [switch   mib2   brief   packets   detail   errors] [nonzero]
To display the causes configured to place operating status to a down or dormant state for one or more ports:	<b>show port operstatuscause</b> [admin   any   cos   dot1x   flowlimit   init   lag   linkflap   linkloss   modifiable   policy   self] [port-string]

**Table 6-3 Displaying Port Configuration Information and Statistics (continued)**

Task	Command
To display all ingress-filter enabled ports or the ingress-filter state of the specified ports:	<b>show port ingress-filter</b> <i>port-string</i>
To display alias name(s) assigned to one or more ports:	<b>show port alias</b> [ <i>port-string</i> ]
To display the status of the force link down function:	<b>show forcelinkdown</b>
To display port transceiver information:	<b>show port transceiver</b> [ <i>port-string</i> ] [ <b>basic-only</b> ] [ <b>sensor-only</b> ] [ <b>all</b> ]
To display the default speed setting on one or more ports:	<b>show port speed</b> [ <i>port-string</i> ]
To display the default duplex setting for one or more ports:	<b>show port duplex</b> [ <i>port-string</i> ]
To display the status of jumbo frame support and MTUs on one or more ports:	<b>show port jumbo</b> [ <i>port-string</i> ]
To display the status of auto-negotiation for one or more ports:	<b>show port negotiation</b> [ <i>port-string</i> ]
To display MDIX mode on one or more ports:	<b>show port mdix</b> [ <i>port-string</i> ] { <b>all</b>   <b>auto</b>   <b>mdi</b>   <b>mdix</b> }
To display the advertised abilities on one or more ports:	<b>show port advertise</b> [ <i>port-string</i> ]
To display the flow control state for one or more ports:	<b>show port flowcontrol</b> [ <i>port-string</i> ]
To display the default 802.1D priority for one or more ports:	<b>show port priority</b> [ <i>port-string</i> ]
To display port broadcast suppression information on one or more ports:	<b>show port broadcast</b> [ <i>port-string</i> ]
To display transmit queue monitoring configuration settings:	<b>show txqmonitor settings</b>
To display transmit queue monitoring port values:	<b>show txqmonitor port</b>
To display port operstatus cause status:	<b>show port operstatuscause txqmn</b> [ <i>port-name</i> ]

## Terms and Definitions

[Table 6-4](#) lists terms and definitions used in this port configuration discussion.

**Table 6-4 Port Configuration Terms and Definitions**

Term	Definition
auto-negotiation	An Ethernet feature that facilitates the selection of port speed, duplex, and flow control between the link segments by first advertising these capabilities and then selecting the fastest transmission mode common to both segments.
baud rate	The speed the console port operates at.
broadcast suppression	A port feature that sets a threshold on the broadcast traffic that is received and switched out to other ports.

**Table 6-4 Port Configuration Terms and Definitions (continued)**

<b>Term</b>	<b>Definition</b>
console port	A port through which local management of the device can be accessed using a terminal or modem.
default priority	A default 802.1p priority that will be applied to a packet when no priority is set in the packet as it transits the port.
duplex	The specification of whether the communications between two devices is one way at a time or both ways simultaneously.
flow control	A port feature that manages the rate of data transmission between two nodes to prevent a fast sender from overrunning a slow receiver.
force linkdown	A port feature that allows for the forcing of a port in the “operstatus down” state to become disabled.
ingress filtering	A port feature that provides a means of limiting the forwarding of received frames on the ingress port based on the VLAN egress list for that port.
jumbo frame	A port feature that supports Ethernet frames greater than 1500 bytes of payload on the port.
link flap detection	A port feature that detects when a link is rapidly going up and down and provides for a port behavior when a threshold is crossed during a configured interval.
MDI/MDIX	A port feature that detects and adapts to straight through (MDI) or cross-over (MDIX) Ethernet cabling on the switch ports.
port advertised ability	The aspect of auto-negotiation that allows a port to share its capabilities with the other end of the link.
port alias	The association of a string name with a port.
port string	A port identifier made up of port type, chassis slot the module containing the port is installed into, and the port number, delineated by a period (.).

## ***Ethernet Operations, Administration, and Maintenance (OAM) Configuration***

This document provides the following information about configuring Ethernet Operations, Administration, and Maintenance (OAM) on the 7100-Series platforms.

<b>For information about...</b>	<b>Refer to page...</b>
<a href="#">Using Ethernet OAM in Your Network</a>	7-1
<a href="#">Implementing Ethernet OAM</a>	7-2
<a href="#">Ethernet OAM Overview</a>	7-2
<a href="#">Configuring Ethernet OAM</a>	7-8
<a href="#">Ethernet OAM Configuration Example</a>	7-10
<a href="#">Terms and Definitions</a>	7-11

### **Using Ethernet OAM in Your Network**

Ethernet Operations, Administration, and Maintenance (OAM) is a collection of standards provided by multiple standards bodies to enable network operators a means to effectively monitor and troubleshoot individual Ethernet links. The Extreme Networks modular switch OAM implementation supports the IEEE 802.3-2008 Clause 57 standard.

The IEEE 802.3-2008 Clause 57 standard allows network operators to monitor and exercise an individual Ethernet link. It provides a set of diagnostics and monitoring functions at a data link level, allowing operators to make a determination of a link's relative health and operational status, and to take administrative action against degraded or faulty links.

A network operator may use SNMP to periodically poll devices for statistics in an attempt to determine when faults occur. However, when a fault occurs, the network operator does not have any means of detecting which links are impacted, and which customers are affected. By implementing OAM across the network, the network operator can proactively determine link degradation or failure, and indicate which customer services are down.

An operator may choose to define an administrative action that will take effect when an error condition occurs or a specified error threshold is crossed. This action may include the generation of Syslog events to bring the state of the suspect link to the operator's attention or to operationally disable the link to avoid network service interruptions caused by degraded or faulty links.

Once corrective action has taken place, link monitoring can be used to verify that the remedial action has succeeded.

The OAM implementation includes a Unidirectional Link Detection (ULD) feature capable of determining when an otherwise bidirectional link is only operational in one direction. This

problem is often caused by faulty wiring or a hardware failure. ULD is not explicitly defined in the OAM standard. The OAM standard provides the means for a ULD solution through a combination of its discovery protocol, administrative actions, and organization specific information LLDP Type-Length-Value (TLV)s. ULD provides for disabling the offending port and sending a Syslog message or only sending a Syslog message.

## Implementing Ethernet OAM

To implement Ethernet OAM on your network:

1. Set the port OAM status to enabled for ports to be monitored.
2. Optionally, set the port OAM mode to passive for any ports that should not initiate OAM contact with their neighbors.
3. Optionally, modify the default OAM link monitor threshold configuration and actions to be taken when link monitor thresholds are crossed, for each OAM enabled port.
4. Optionally, configure the number of event notification OAM PDUs retransmitted by the remote peer.
5. If you need to bring up an interface that has been taken down due to exceeding OAM monitoring thresholds, clear the OAM operstatus cause to bring the interface back up.
6. Optionally, configure ports for ULD.

## Ethernet OAM Overview

For information about...	Refer to page...
<a href="#">OAM Client</a>	<a href="#">7-2</a>
<a href="#">OAM Discovery</a>	<a href="#">7-2</a>
<a href="#">OAM Client Mode</a>	<a href="#">7-3</a>
<a href="#">OAM Datalink Layer Monitoring</a>	<a href="#">7-3</a>
<a href="#">OAM Event Notification Retries</a>	<a href="#">7-7</a>

### OAM Client

The OAM client contains the essential control operations and state information concerning OAM operations on a specific port. It is responsible for the handling of received OAM PDUs from remote clients, and based upon the state of local and remote settings, allows OAM to operate upon a link.

Link events are transmitted via OAM PDUs between OAM client entities. The OAM Client is also responsible for maintaining statistics concerning transmitted and received OAM PDUs.

You must enable the OAM client on the device for OAM operations to take place. The OAM client is disabled by default.

Use the **set port oam status** command to enable the OAM client on the device.

### OAM Discovery

Periodic Information OAM PDU messages are exchanged between OAM clients to both initiate OAM discovery on the link and, once initiated, assure that remote client information is correct

from the perspective of the local client. Information OAM PDUs can contain the remote OAM client's information, as well as a copy of the local client's information. Both clients must accept the exchanged information to complete OAM discovery. Clients can reject received information that is incorrect, outdated, or incomplete.

OAM clients must accept two sets of information to complete OAM discovery:

- The remote client's information
- A copy of the local client's information that has been reproduced by the remote client

Once discovery is completed, any change to configuration or state information, on either OAM client, forces the OAM discovery session to be torn down and re-established with the new information.

## OAM Client Mode

OAM clients may operate in either active or passive mode. Clients configured for active mode may initiate contact with remote peers.

OAM clients configured for passive mode may not initiate any contact with a remote peer. Passive OAM clients are only allowed to respond to requests received from a remote peer. When OAM is disabled on the port, all OAM frames are discarded.

You may want to require that a remote peer outside of your administrative control be set to OAM passive mode to disable that port's ability to initiate OAM discovery and processing, as well as to prevent the remote port from setting the local port in remote loopback.

OAM client mode is configured using the **set port oam mode** command.

## OAM Datalink Layer Monitoring

OAM includes datalink layer monitoring, which can be used to proactively take action upon links which are exhibiting faulty behavior.

The OAM link monitor checks for symbol and frame errors.

- A symbol is a fixed length electrical or optical waveform on the wire that represents a binary value and is a subcomponent of a frame. A symbol error is a symbol with an invalid data value.
- A frame is made up of a variable number of symbols formatted with a preamble, header, data blocks, and checksum. A frame error can have a number of causes such as MTU exceeded, invalid frame check sequence, an alignment error (length is not an integer number of octets) or length error.

Network administrators may define threshold values for symbol or frame errors. If the threshold is exceeded within a configured link monitor window, a threshold crossing event has transpired and an action will be taken. Supported link monitor actions include:

- Transmitting a Link Event Notification message to the remote OAM client
- Generating a Syslog message
- Taking the link operationally off-line

The administrator has the option of keeping a suspect link operationally disabled until manual intervention has been taken or for the link to be operationally re-enabled once the error condition has resolved.

Configure OAM datalink layer monitoring using the **set port oam link-monitor** command.

Configuring OAM datalink layer monitoring includes setting several options, described in the following sections.

## Frame Option

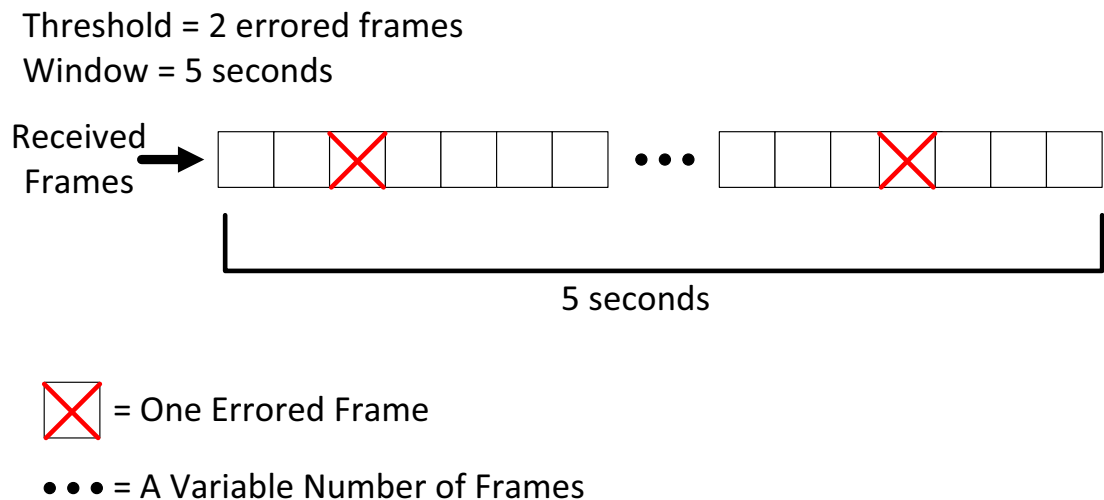
The **frame** option monitors frame errors occurring during a period of time. The default threshold for the frame option is **1** errored-frame. The default window for the frame option is **1** second, and the maximum window is **60** seconds.

As presented in [Figure 7-1](#):

- The configured threshold is **2** errored-frames that occur within the configured window.
- The window within which the threshold must be exceeded for a configured action to occur is **5** seconds.

If more than **2** errored-frames are received on the monitored port within any **5** second window, the configured action occurs.

**Figure 7-1 Frame Link Monitor Option**



## Frame-Seconds Option

The **frame-seconds** option monitors frames within one or more one-second windows for errors. The link monitor window is specified as a number of one-second windows to monitor. The default threshold is **1** errored-second window. The default window is **60** one-second windows. The minimum window is **10** one-second windows, and the maximum window is **900** one-second windows.

For example, if defaults are being used and a single frame error occurs within one of sixty one-second windows, the threshold has been met. If the threshold was raised to three errored-seconds, three one-second windows would have to register a frame error, within sixty one-second windows, for the threshold to be met.

As presented in [Figure 7-2](#):


- The configured threshold is **2** one-second windows in which one or more errored-frames occur.
- The window is **25** one-second windows to monitor.

If there is one or more frame errors in 3 or more one-second windows within the link monitor window of 25 one-second windows, the configured action occurs.

**Figure 7-2 Frame-Seconds Link Monitor Option**

Threshold = 2 errored one-second windows

Window = 25 one-second windows



### Frame-Period Option

The **frame-period** option monitors frame errors that occur during the reception of a given number of frames. The default threshold is 1 errored-frame. The default window is equivalent to the maximum number of minimum sized frames that may be transmitted over the link during a 1 second interval, and the upper bound is the maximum number of minimum sized frames that may be transmitted over the link during a 1 minute interval.

The frame-period option defines its window values based upon the line rate of the port being configured. As such, option values may not be determined until the port has achieved a valid link state. The frame-period window default and range values, based upon link speed, are displayed in [Table 7-1](#).

**Table 7-1 Frame-Period Window Values**

Port Line Rate	Default Window Value	Window Value Range
100 Mbps	148,800 frames	14880 - 8928000 frames
1Gbps	1,488,000 frames	148,800 to 89,280,000 frames
10 Gbps	14,880,000 frames	1,488,000 to 892,800,000 frames

As presented in [Figure 7-3](#):

- The configured threshold is 4 errored-frames that occur within the configured window.
- The window is the default window for a 1Gbps port: 1,488,000 frames.


If there is more than 4 errored-frames within the reception of 1,488,000 frames, the configured action occurs.

**Figure 7-3 Frame-Period Link Monitor Option**




Threshold = 4 Errored Frames

Window = 1,488,000 Frames



••• = 1487971 Frames

 = One Errored Frame

### Symbol-Period Option

The **symbol-period** option monitors symbol errors that occur during the reception of a given number of symbols. The default threshold is **1** errored-symbol. The default window is equivalent to the maximum number of symbols that may be transmitted over the link during a one second interval, and the upper bound is the number of symbols that may be transmitted over the link during a one minute interval.

The symbol-period option defines its default window value based upon the line rate of the port being configured. As such, the default window value may not be determined until the port has achieved a valid link state. The symbol-period window default and range values, based upon link speed, are displayed in [Table 7-2](#) on page 7-6.

**Table 7-2 Symbol-Period Window Values**

Port Line Rate	Default Window Value	Window Value Range
100 Mbps	131,072,000 symbols	52,428,800 - 7,864,320,000
1Gbps	524,288,000 symbols	524,288,000 - 31,457,280,000
10 Gbps	5,242,880,000 symbols	5,242,880,000 - 314,572,800,000

As presented in [Figure 7-4](#):


- The configured threshold is **5** errored-symbols.
- The window is the default window for a 1Gbps port: **524,288,000** symbols.

If there is more than **5** errored-symbols within the reception of **524,288,000** symbols, the configured action occurs.

Figure 7-4 Symbol-Period Link Monitor Option

Threshold = 5 Errored Symbols

Window = 524,288,000 Symbols



••• = 524,287,971 Symbols

⊗ = One Errored Symbol

### Actions

The administrator may configure one or more of three actions to be taken upon the detection of a link event:

- The **syslog** option triggers a syslog message to be generated, which confers information related to the event.
- The **notify** option triggers the transmission of a link event notification OAM PDU to the remote client.
- The **disable-interface** option operationally disables the port in question, and the port remains in that state until the administrator restores the port to operational status.

### OAM Event Notification Retries

The local OAM client can be configured to retransmit event notification OAM PDUs to its remote peer up to the configured number of retries.

If an OAM monitor link threshold is crossed, generating an event, and the notification action is configured, a single event notification OAM PDU is always generated. By default no further notifications are sent unless a value greater than zero is configured for the number of OAM event retries. This retransmission process will halt if the link monitoring process determines that additional events have transpired upon the link. A notification of the new event is then sent by the remote OAM client. The default value is **0**, and the maximum value is **10**.

Use the **set port oam notify-retry** command to configure the number of event notification retries sent by the remote OAM client when an OAM event occurs.

### Unidirectional Link Detection (ULD)

When a link is in a unidirectional state, the ability to pass traffic over the link is broken in one direction. This causes problems for protocols that depend upon reliable bidirectional communications to operate. For example, the Spanning Tree Protocol (STP) relies on the bidirectional exchange of bridge protocol data units to create a loop free topology. By implementing link layer OAM with Unidirectional Link Detection (ULD) across the network, OAM ULD can be configured to take a port out of service when a unidirectional link is detected. In the case of a unidirectional link in a STP context, the creation of a loop within the network is prevented.

## ULD Mode

By default ULD mode is disabled, and the detection of unidirectional links does not occur. Extreme Networks supports two means of detecting a unidirectional link on a port using ULD modes:

- **Standard** – Uses the existing OAM discovery protocol to perform the configured ULD action, if more than 5 seconds elapses between reception of standard information OAMPDUs on the port.
- **Fast** – Establishes a second tier of OAM discovery by transmitting information OAMPDUs with the Fast ULD information TLV. The configured ULD action is performed when up to 3 times the interval defined by the fast timer setting (a configurable range of between 600 milliseconds and 3 seconds) elapses between reception of a Fast ULD information TLV on the port.

You can set the ULD mode on a per port basis using the **set port oam uld mode** command.

## ULD Fast Timer

The ULD fast timer determines the interval between transmission of Fast ULD information TLVs used by the ULD fast mode to detect unidirectional links.

When setting the fast timer you specify a fast timer multiplier between **2** and **10**. ULD sets the fast timer value by multiplying the specified fast timer multiplier by 100 milliseconds. The fast timer interval defaults to 200 milliseconds (2 x 100 milliseconds). When using the default timer interval of 200 milliseconds, up to 600 milliseconds (3 times the configured timer interval) will elapse between the reception of the Fast ULD information TLV and ULD performing the configured ULD action.

You configure the fast timer setting by specifying the fast timer multiplier using the **set port oam uld fast-timer** command.

## ULD Action

When ULD detects a unidirectional link in either standard or fast mode, a Syslog message is sent or a Syslog message is sent and the port is disabled. The default action is sending a Syslog message without disabling the port.

You configure the ULD action on a port when a unidirectional link is detected using the **set port oam uld action** command.

# Configuring Ethernet OAM

This section provides:

- A listing of OAM default values
- An OAM configuration procedure

[Table 7-3](#) lists the 7100-Series device default Ethernet OAM configuration settings.

**Table 7-3 Default Ethernet OAM Configuration Settings**

Parameter	Description	Default Value
OAM client mode	Specifies the operating mode of the local OAM client for the port.	active
OAM event notification retries	Sets the number of event notification retries to send for the port.	0

**Table 7-3 Default Ethernet OAM Configuration Settings (continued)**

Parameter	Description	Default Value
OAM status	Specifies the state of the local port OAM client.	disabled
OAM monitor link frame window	Specifies the period for single errored-frame threshold.	1 second
OAM monitor link frame-period frame threshold	Specifies the number of errors received based upon the number of frames per second by port linked rate.	100 Mbps – 1 error per 148,800 frames per second 1Gbps – 1 error per 1,488,000 frames per second 10 Gbps – 1 error per 14,880,000 frames per second
OAM monitor link frame-seconds threshold	Specifies the number of one-second intervals in which one or more frame errors occurs.	1 errored-second
OAM monitor link symbol-period threshold	Specifies the number of symbol errors that occur during the reception of a given number of symbols based upon the port line rate.	100 Mbps – 1 error per 131,072,000 symbols per second 1Gbps – 1 error per 524,288,000 symbols per second 10 Gbps – 1 error per 5,242,880,000 symbols per second
OAM monitor link action	Specifies the action that occurs should a link monitor event occur.	notify
OAM ULD mode	Specifies whether ULD is using standard or fast mode to detect unidirectional links or whether ULD is disabled.	disabled
OAM ULD action	Specifies whether ULD will only send a Syslog message or both disable the port and send a Syslog message when a unidirectional link is detected.	send Syslog message only
OAM ULD fast timer multiplier	Specifies an integer value, multiplied by 100 milliseconds, to determine the fast timer interval used by ULD fast mode for sending OAMPDUs with Fast ULD information TLVs.	2

[Procedure 7-1](#) provides an example of an OAM configuration.

#### Procedure 7-1 Configuring OAM

Step	Task	Command(s)
1.	Administratively enable Ethernet OAM on the specified port(s).	<b>set port oam <i>port-string</i> status enable</b>
2.	Optionally, set the operating mode for the OAM client on the specified port(s).	<b>set port oam <i>port-string</i> mode {active   passive}</b>

**Procedure 7-1 Configuring OAM**

Step	Task	Command(s)
3.	Optionally, configure OAM link monitor functionality for the specified port.	<b>set port oam <i>port-string</i> link-monitor {frame   frame-period   frame-seconds   symbol-period} {threshold <i>threshold</i>   window <i>window</i>   action {[syslog] [disable-interface] [notify]}</b>
4.	Optionally, set the number of notify retries to send for the specified port.	<b>set port oam <i>port-string</i> notify-retry <i>retries</i></b>
5.	Optionally, bring up an interface that has been taken down due to exceeding OAM monitoring thresholds.	<b>clear port operstatuscause [<i>port-string</i>] [oam]</b>
6.	Optionally, enable ULD standard or fast mode to detect unidirectional links on a port.	<b>set port oam <i>port-string</i> uld mode {standard   fast}</b>
7.	Optionally, set the ULD action to both disable the unidirectional port and send a Syslog message or only send a Syslog message.	<b>set port oam <i>port-string</i> uld action {disable-port   syslog-only}</b>
8.	Optionally, change the ULD fast timer multiplier for a port.	<b>set port oam <i>port-string</i> uld fast-timer <i>multiplier</i></b>

## Ethernet OAM Configuration Example

The following CLI example provides a local and remote port OAM configuration that:

- Enables port OAM status on both the local (ge.3.14) and remote (ge.1.10) ports.
- Retains the default active OAM mode on the local port and sets the remote port OAM mode to passive. This configuration allows the local port to initiate OAM on the link. This configuration prevents the remote port from initiating any OAM operations.
- Sets the number of notify retries sent by the remote port to 3. The local port retains the default value of 0 retries.
- Sets the local and remote port's OAM monitor link frame window to 100 seconds, the threshold to 2 errors and the action to Syslog, disable the interface, and notify, should the frame event threshold be reached.
- Enables OAM ULD standard mode on all ports, both local and remote.
- Sets the ULD action to both disable the port and send a Syslog message on all ports, both local and remote.

**Switch 1, Local Port: ge.3.14**

```
7100 Chassis(rw)->set port oam ge.3.14 status enable
7100 Chassis(rw)->set port oam ge.3.14 notify-retry 3
7100 Chassis(rw)->set port oam ge.3.14 link-monitor frame threshold 2
7100 Chassis(rw)->set port oam ge.3.14 link-monitor frame window 100
7100 Chassis(rw)->set port oam ge.3.14 link-monitor frame action syslog
disable-interface notify
7100 Chassis(rw)->set port oam *.*.* uld mode standard
7100 Chassis(rw)->set port oam *.*.* uld action port-disable
```

## Switch 2, Remote Port ge.1.10

```

7100 Chassis(rw)->set port oam ge.1.10 status enable
7100 Chassis(rw)->set port oam ge.1.10 mode passive
7100 Chassis(rw)->set port oam ge.1.10 link-monitor frame threshold 2
7100 Chassis(rw)->set port oam ge.1.10 link-monitor frame window 100
7100 Chassis(rw)->set port oam ge.1.10 link-monitor frame action syslog
disable-interface notify
7100 Chassis(rw)->set port oam *.*.* uld mode standard
7100 Chassis(rw)->set port oam *.*.* uld action port-disable

```

## Terms and Definitions

Table 7-4 lists terms and definitions used in this OAM configuration discussion.

**Table 7-4 OAM Configuration Terms and Definitions**

Term	Definition
frame	A frame is made up of a variable number of symbols formatted with a preamble, header, data blocks, and checksum.
frame monitoring	Monitors link frame errors that occur during a single specified window in seconds.
frame-period monitoring	Monitors frame errors that occur during the reception of a given number of frames relative to the port line rate.
frame-seconds monitoring	Monitors a threshold made up of the number of one second intervals in which one or more frame errors occurs over a specified window of one second intervals.
IEEE 802.3-2008 Clause 57	An OAM standard that allows network operators to monitor and exercise an individual Ethernet link.
OAM	Ethernet Operations, Administration, and Maintenance (OAM) is a collection of standards provided by multiple standards bodies to enable network operators a means to effectively monitor and troubleshoot individual Ethernet links.
OAM client	Contains the essential control operation and state information concerning OAM operations on a specific port, including the handling of received OAM PDUs from the remote client.
OAM client mode	Specifies whether an OAM client can initiate contact with remote peers.
OAM datalink layer monitoring	Monitors Ethernet links for faulty link behavior and is capable of proactive actions when error notifications are received from a remote OAM client.
symbol	A fixed length electrical or optical waveform on the wire that represents a binary value and is a subcomponent of a frame.
symbol-period monitoring	Monitors the number of symbol errors that occur during the reception of a given number of symbols relative to the port line rate.
Unidirectional Link Detection (ULD)	An OAM port feature capable of detecting a link that is only operational in one direction.
ULD action	The action that occurs when a unidirectional link is detected to both disable the port and send a Syslog message or only send a Syslog message.
ULD fast mode	A ULD mode of operation that establishes a second tier of OAM discovery by <ul style="list-style-type: none"> <li>Transmitting information OAMPDUs with the Fast ULD information TLV</li> <li>Performing the configured ULD action if up to 3 times the fast timer period elapses between reception of Fast ULD information TLVs on the port.</li> </ul>

**Table 7-4 OAM Configuration Terms and Definitions (continued)**

Term	Definition
ULD fast timer multiplier	An integer value that is multiplied by 100 milliseconds to determine the ULD fast timer interval used by ULD fast mode for transmission of OAMPDUs with Fast ULD information TLVs.
ULD standard mode	A ULD mode of operation that uses the existing OAM discovery protocol to perform the configured ULD action on the port, if more than 5 seconds elapses between reception of standard information OAMPDUs on the port.

## Port Mirroring Configuration

This chapter provides the following information about configuring and monitoring port mirroring on 7100-Series devices.

For information about...	Refer to page...
<a href="#">How to Use Port Mirroring in Your Network</a>	8-1
<a href="#">Implementing Port Mirroring</a>	8-3
<a href="#">Overview of Port Mirroring Configurations</a>	8-3

### How to Use Port Mirroring in Your Network

Port mirroring, also known as port redirect, is a network traffic monitoring method. It forwards a copy of each received or transmitted frame (or both) from one or more switch ports (source ports) to another port or ports (destination ports) where the data can be studied. Once the bit stream from one or more source ports is mirrored to one or more destination ports, you can further analyze the captured data using an RMON probe or a network sniffer without affecting the original port's normal switch operation.

Port mirroring is an integrated diagnostic tool for tracking network performance and security that is especially useful for fending off network intrusion and attacks. It is a low-cost alternative to network taps and other solutions that may require additional hardware, may disrupt normal network operation, may affect client applications, and may even introduce a new point of failure into your network. Port mirroring scales better than some alternatives and is easier to monitor. It is convenient to use in networks where ports are scarce.

Mirroring egress traffic results in the mirrored traffic always having an 802.1Q VLAN tag. The VLAN and priority values are those that were used for transmission of the original packet.

You can set up the following types of port mirroring relationships on received or transmitted traffic (or both):


- One-to-one (source port to destination port)
- Many-to-one
- One-to-many

Depending on your network, ports that you can configure to participate in mirroring include physical ports and host ports—VLAN interfaces. For more information, refer to “[Overview of Port Mirroring Configurations](#)” on page 8-3.

You can use port mirroring for analyzing bi-directional traffic and ensuring connectivity between, for example, a departmental switch and its high speed uplink to your backbone switch as shown in [Figure 8-1](#).




**Figure 8-1 Using Port Mirroring to Monitor a Departmental Switch**



This one-to-one configuration would allow you to capture traffic in both directions to the backbone uplink port. In this example, you would set a port mirror between departmental switch port 4.1 (source) and the destination port 4.2 connected to the traffic probe.

You can also use port mirroring, for example, to monitor all received traffic or a specific type of received traffic to your backbone switch as shown in [Figure 8-2](#).

**Figure 8-2 Using Port Mirroring to Monitor Incoming Traffic to a Backbone Switch**



The many-to-one configuration in this example would be possible by setting a port mirror on the backbone between source ports 1.2 and 2.1 to destination port 1.1.

The 7100-Series platform supports a maximum of two unique destinations in the following configurations:

- Up to two one-to-one mirrors
- Up to two many-to-one mirrors
- One one-to-two mirror

This allows configurations like:

- a. Up to two one-to-one mirrors
- b. Up to two many-to-one mirrors, or
- c. A single one-to-two mirror.

For the “one-to-many” there can be up to 2 destination ports.

For the “many-to-one” there is no limit to the number of source ports.

For the port mirror case the source port(s) can be a physical port or VLAN.

LAG ports can not be used as the source port for a mirror.

Mirror destinations can be physical ports or LAGs, including ones on other switches in the same stack. Mirror destinations can not be VLANs.

The port and VLAN mirror function does not mirror error frames.

Mirroring egress traffic results in the mirrored traffic always having an 802.1Q VLAN tag. The VLAN and priority values are the ones used for transmission of the original packet.

## Implementing Port Mirroring

You can implement port mirroring on 7100-Series devices using simple CLI commands. The source port of a VLAN mirror is a VTAP interface created using the **set vlan interface** command. A VTAP interface provides the data source input of a VLAN mirror and must exist before attempting to create a VLAN port mirror. Once the specific device ports are operationally linked, use the **set port mirroring** command to create a mirroring relationship between your intended source and your destination ports.

You can also use CLI to operationally disable mirroring, if necessary, and to specify whether to mirror received traffic, transmitted traffic, or both. You can also monitor multicast traffic by enabling IGMP mirroring on specific ports.



**Note:** It is important to not oversubscribe ports in a mirroring configuration. This can cause bottlenecks and will result in discarded traffic.

Once configured, all packets (network, data, control, and so on) received by the switch will be mirrored. Errored packets will not be mirrored. Unless you disable Spanning Tree on destination ports, they will continue to function as active bridge ports, in accordance with the SMON (Switch Monitoring) standard.

## Overview of Port Mirroring Configurations

One or more source ports can be mirrored locally to another physical port within the same 7100-Series device. In addition, virtual ports and other types of port configurations can also participate in mirroring on Extreme Networks switching devices as described in the following sections:

- [LAG Mirrors](#)
- [VLAN Mirrors](#)

### LAG Mirrors

Each 7100-Series module designates a specific number of virtual link aggregation ports which the Link Aggregation Control Protocol (LACP) can use to dynamically group multiple physical ports into one logical link. Once underlying physical ports (such as ge.x.x) are associated with an

aggregator port, the resulting aggregation is represented as one Link Aggregation Group (LAG) with a lag.x.x port designation.

Refer to the [Chapter 21, Link Aggregation Control Protocol \(LACP\) Configuration](#) for more information.

The 7100-Series only supports LAGs as mirror destinations. A LAG is not supported as a mirror source.

## VLAN Mirrors

Creating a VLAN and setting a mirror for the VLAN allows you to monitor all traffic to your specified VLAN interface. For example, you could track all data traveling in and out of a confidential group of workstations, such as a Finance VLAN, by analyzing only one connection point. Considerations when configuring VLAN mirrors include:

- A one-to-many or many-to-one VLAN mirror is considered a single destination port.
- Many-to-one mapping allows multiple VLANs to be sent to one specific destination port.
- Oversubscribed traffic will be dropped.

A VTAP interface provides the data source input of a VLAN mirror. VTAP creation is the mechanism for adding a MIB-II interface table entry for a VLAN. A VLAN will not have a MIB-II ifIndex if a VTAP interface does not exist for it. Use the **set vlan interface** command to create a VTAP interface.

## Avoiding Bottlenecks

It is especially important to not oversubscribe ports in a mirroring configuration because this can cause bottlenecks and will result in discarded traffic.

If, for example, there are 10 users in VLAN 1, each attached to a 10 Mbps port, when you mirrored VLAN 1 to another 10 Mbps port to which your sniffer is attached, the probe switch would probably have to drop packets at the destination port. Since your purpose in configuring mirroring is to see all of the traffic for VLAN 1, it would be better in this scenario to attach the sniffer to a 100 Mbps port.

## Configuring Port Mirrors



**Note:** When a port mirror or policy mirror destination is created, It is automatically enabled.

For information about...	Refer to page...
<a href="#">Reviewing Port Mirroring</a>	8-4
<a href="#">Setting Port or VLAN Mirroring</a>	8-5
<a href="#">Deleting Mirrors</a>	8-6

## Reviewing Port Mirroring

Use this command to display the status of port mirroring and information about any mirrors configured:

```
show port mirroring
```

## Examples

This example shows that no port mirrors are configured on the device:

```
7100 Chassis(rw)->show port mirroring
No Port Mirrors configured.
IGMP Multicast Mirror status Disabled
Mirror Outbound Rate Limited Frames : Disabled
```

This example shows that a port mirror is configured between source port vtap.0.5 and ge.1.1 and that both received (Rx) and transmitted (Tx) frames will be monitored. It also shows that mirroring status is currently administratively and operationally enabled. A mirror must be administratively enabled (as described in the next section) and its source and destination ports must have an active link for operational status to be enabled.

```
7100 Chassis(rw)->show port mirroring
Port Mirroring
=====

Source Port      = vtap.0.5
Target Port      = ge.1.1
Frames Mirrored  = Rx and Tx
Admin Status     = enabled
Operational Status = enabled

Mirror Outbound Rate Limited Frames : Disabled
```

## Setting Port or VLAN Mirroring

Use this command to create a new mirroring relationship, or to enable or disable an existing mirroring relationship. Optionally, you can specify whether to mirror received frames, transmitted frames, or both:

```
set port mirroring {create | disable | enable} source destination [both | rx
| tx]
```

If not specified, **both** received and transmitted frames will be mirrored.



**Note:** By default, when you create a port mirror, the port mirror is enabled.

## Examples

This example shows how to create a port mirror to mirror frames sourced on port ge.1.4 and received on port ge.1.11:

```
7100 Chassis(rw)->set port mirroring create ge.1.4 ge.1.11 rx
```

This example shows how to create a many-to-one mirroring configuration between source ports ge.1.2, ge.1.3 and ge.1.4, and target port ge.1.10.:

```
7100 Chassis(rw)->set port mirroring create ge.1.2-4 ge.1.10
```

This example shows how to configure mirroring from source port 5 to destination port 1 in slot 1 (ge.1.1):

```
7100 Chassis(rw)->set vlan interface 5 create
7100 Chassis(rw)->set port mirroring create vtap.0.5 ge.1.1
7100 Chassis(rw)->show port mirror
Port Mirroring
=====

Source Port      = vtap.0.5
Target Port      = ge.1.1
Frames Mirrored  = Rx and Tx
Admin Status     = enabled
Operational Status = Unavailable resources

Mirror Outbound Rate Limited Frames : Disabled
```



**Note:** If you configure a port mirror on an uplink (tagged) port, make sure the port is assigned to egress frames with that VLAN tag. Refer to [Chapter 24, VLAN Configuration](#) for more information about configuring VLANs.

## Deleting Mirrors

Use this command to clear a port mirroring configuration:

```
clear port mirroring source destination
```

## Remote Mirroring Using a Layer 2 GRE Tunnel

The 7100-Series supports remote mirroring using a Layer 2 (L2) GRE tunnel. Any L2 traffic (unicast, multicast, or broadcast) that can be mirrored with a non-L2 GRE remote mirror can be mirrored with an L2 GRE remote mirror. The mirror source port is the source of the mirrored packets found on the local router of interest. The mirror encapsulates the L2 traffic seen by the mirrored source port in an IP GRE header and delivers it to the L2 tunnel bridge port as the destination address. The tunnel destination address is the ultimate destination port of the tunnel where packets are decapsulated and delivered to the port local to the remote router.



**Note:** The 7100-Series supports encapsulation for remote mirror; the 7100-Series does not support decapsulation. The remote mirror destination device must be a device that supports decapsulation for full remote mirroring support. An 7100-Series can be used if you intend to use a packet monitoring program as the means of capturing the packet contents.

The 7100-Series supports IPv4 addressing only for remote mirroring.

The tunnel destination port resides on a remote router with the appropriate L2 GRE tunnel configuration back to the ultimate mirror destination. The intermediate mirror destination port is the actual mirror destination of the local SMON. Once the L2 tunnel is enabled and fully configured, packets sent to an intermediate destination are encapsulated and forwarded for any SMON port mirrors that use the mirror destination port. Once the tunnel is enabled, the tunnel destination port is in internal loopback mode and can no longer be used as a normal switch port.

Configuration of a remote mirror using a L2 GRE tunnel consists of:

- Creating a tunnel interface in global configuration mode, with a loopback source address and a route to the destination
- Configuring the mirrored port:
  - Configure an SMON port mirror by creating the port mirror, specifying the L2 GRE tunnel destination as the destination port (must be a tunnel bridge port (**tbp.x.y**))

The following CLI input sets up the L2 GRE tunnel for mirrored port **ge.1.1** by:

- Configuring **loop.0.1** as the loopback source address (**88.88.88.1/32**) for the mirrored port
- Entering configuration mode for tunnel 5
  - Configuring the tunnel 5 destination address (**99.99.99.1**)
  - Setting the tunnel mode to GRE L2 and specifying **tbp.0.1** as the mirror destination
  - Enabling the mirrored tunnel which allocates the necessary resources to support mirrored packets
- Configuring a static route to the mirror destination

```
7100 Chassis(rw) ->configure
7100 Chassis(rw-config) ->interface loop.0.1
7100 Chassis(rw-config-intf-loop.0.1) ->ip address 88.88.88.1/32
7100 Chassis(rw-config-intf-loop.0.1) ->no shutdown
7100 Chassis(rw-config-intf-loop.0.1) ->exit
7100 Chassis(rw-config) ->interface tunnel 5
7100 Chassis(rw-config-intf-tun.0.5) ->tunnel destination 99.99.99.1
7100 Chassis(rw-config-intf-tun.0.5) ->tunnel mode gre l2 tbp.0.1
7100 Chassis(rw-config-intf-tun.0.5) ->tunnel mirror enable
7100 Chassis(rw-config-intf-tun.0.5) ->tunnel source 88.88.88.1
7100 Chassis(rw-config-intf-tun.0.5) ->no shutdown
7100 Chassis(rw-config-intf-tun.0.5) ->exit
7100 Chassis(rw-config) ->ip route 99.99.99.1/32 6.1.1.2 interface vlan.0.20 1
7100 Chassis(rw-config) ->exit
```

The following CLI input creates an SMON port mirror specifying the L2 GRE mirror-destination (**tbp.0.1**) as the destination:

```
7100 Chassis(rw) ->set port mirroring create ge.1.1 tbp.0.1 both
7100 Chassis(rw) ->
```

The 7100-Series does not support decapsulation of the mirrored packet. If you require decapsulation of the mirrored packet, the following CLI input sets up the L2 GRE tunnel on the mirror destination router for an S-Series or K-Series device, by

- Configuring **loop.0.1** as the loopback source address (**99.99.99.1/32**) for the mirrored port
- Configuring VLAN 33 as the VLAN interface the tunnel resides on
- Entering configuration mode for tunnel 12
  - Configuring the tunnel 12 destination address (**88.88.88.1**)
  - Setting the tunnel mode to GRE L2 and specifying **ge.2.4** as the mirror destination

You do not enable the mirrored tunnel on the mirror destination router. The L2 GRE-encapsulated packets:

- Arrive with source **88.88.88.1** and destination **99.99.99.1**

- Are decapsulated and forwarded out physical port **ge.2.4**

A PC running a packet-monitoring program, such as WireShark, can be attached to this port. The packet-monitoring program displays the L2 traffic that is seen by ge.1.1 on the router on which the mirrored port resides.

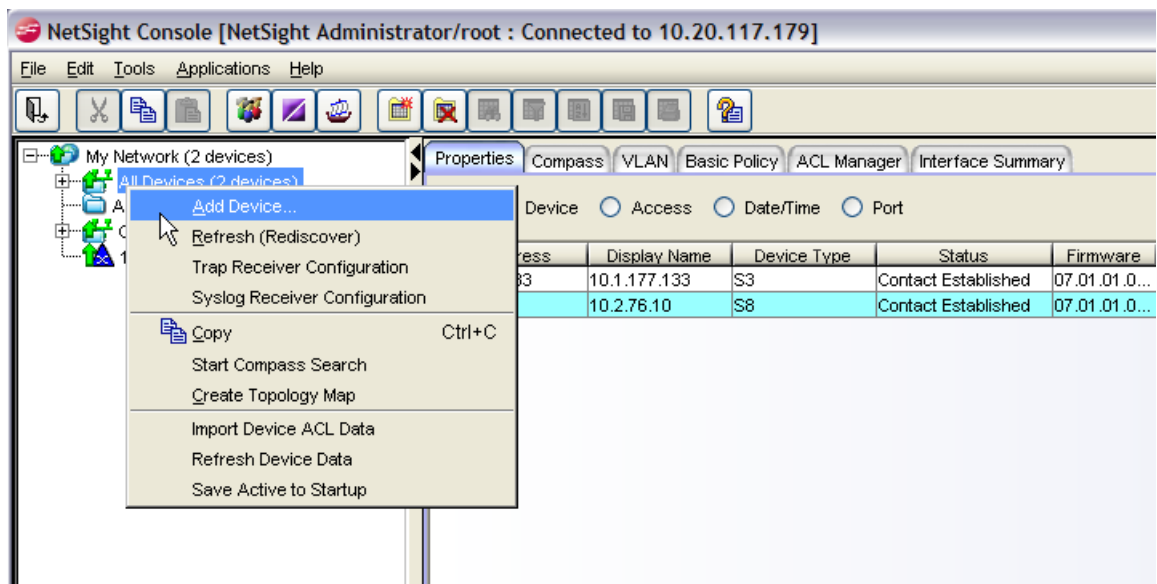
- Configure a static route to the router on which the mirror port resides (**88.88.88.1/32**)

```
S-K-Series(rw)->configure
S-K-Series(rw-config)->interface loop.0.1
S-K-Series(rw-config-intf-loop.0.1)->ip address 99.99.99.1/32
S-K-Series(rw-config-intf-loop.0.1)->no shutdown
S-K-Series(rw-config-intf-loop.0.1)->exit
S-K-Series(rw-config)->interface vlan.0.33
S-K-Series(rw-config-intf-vlan.0.33)->ip address 5.1.1.1 255.255.255.0 primary
S-K-Series(rw-config-intf-vlan.0.33)->no shutdown
S-K-Series(rw-config-intf-vlan.0.33)->exit
S-K-Series(rw-config)->interface tunnel 12
S-K-Series(rw-config-intf-tun.0.12)->tunnel destination 88.88.88.1
S-K-Series(rw-config-intf-tun.0.12)->tunnel mode gre 12 ge.2.4
S-K-Series(rw-config-intf-tun.0.12)->tunnel source 99.99.99.1
S-K-Series(rw-config-intf-tun.0.12)->no shutdown
S-K-Series(rw-config-intf-tun.0.12)->exit
S-K-Series(rw-config)->ip route 88.88.88.1/32 5.1.1.2 interface vlan.0.33 1
S-K-Series(rw-config)->exit
S-K-Series(rw)->
```

## Example: Configuring and Monitoring Port Mirroring

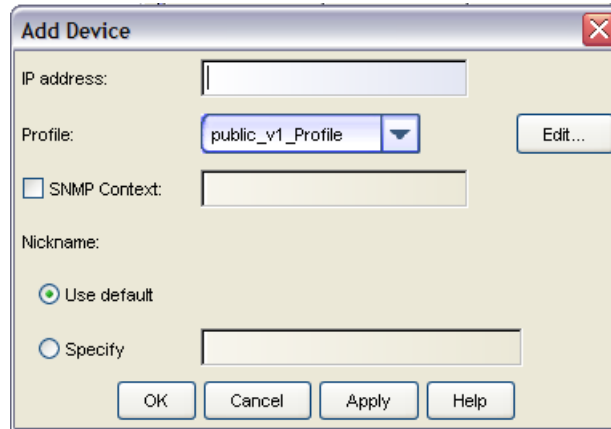
This section describes how to use Extreme Networks NetSight Console from a Network Management Station (NMS) to display RMON statistics for monitoring port mirroring.

1. Log onto Netsight Console.

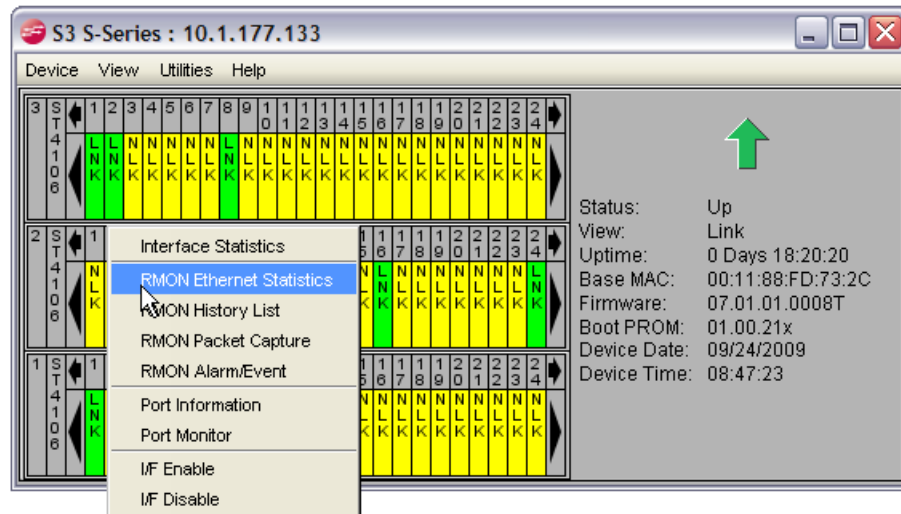


- On the console main screen, expand **My Network** in the file directory tree, right-click **All Devices**, and select **Add Device**.

The Add Device screen displays.



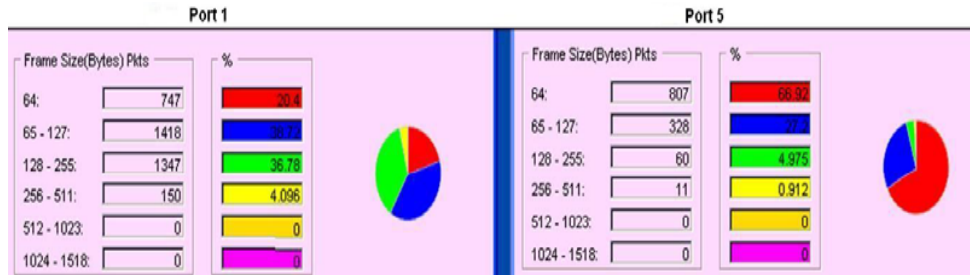
- Model the 7100-Series device by entering its **IP address** in the field provided. Click **OK**.
  - On the console main screen, expand **All Devices** in the file directory tree to show the IP address(es) of the device(s) you just modeled.
  - Right click on the IP address of the 7100-Series device and select **Device Manager**.
- The device manager screen displays for the 7100-Series device.



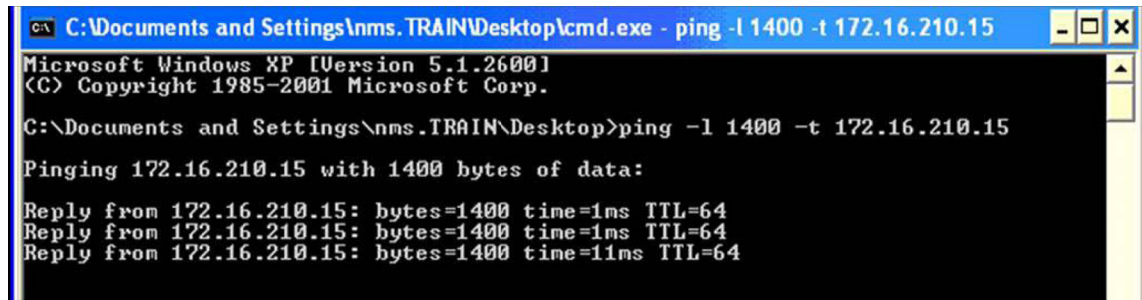
- Right click on port 1 (ge.1.1) and select **RMON Ethernet Statistics**.
- Repeat step 9 for port 5 (ge.1.5).



RMON Ethernet statistics charts will display for ports 1 and 5.



8. Note that the section of the two charts that shows the frame count by frame size lists no larger size frames (512-1518 bytes). In the next step, you will create large frames.
9. Open the Command Prompt window and set up a continuous ping to the 7100-Series device, as shown below. Use `-l 1400` to set the size of the ping frame to 1400 bytes and `-t` to set a continuous ping.



10. Refer back to the RMON Ethernet Statistics windows opened in Steps 9 and 10. You should see the number of **1024 - 1518** frames incrementing on Port 1 because the NMS is connected on this port. You should also see that these larger size frames are not incrementing on Port 5.
11. From the terminal session with the 7100-Series device, create a port mirroring instance with port 1 (ge.1.1) as the source and port 5 (ge.1.5) as the destination port.

```
7100 Chassis(su)->set port mirroring create ge.1.1 ge.1.5 both
```

12. Verify the mirroring configuration.

```
7100 Chassis(su)->show port mirroring
```

```

Port Mirroring
=====
Source Port = ge.1.1
Target Port = ge.1.5
Frames Mirrored = Rx and Tx

Port Mirroring Admin status = enabled
Port Mirroring Oper status = enabled
    
```

13. Refer again to the RMON Ethernet Statistics windows and notice that both port 1 and port 5 are now incrementing the larger size frames. If you connected a network analyzer to port 5, you would see these frames being received and transmitted on port 1.

## System Configuration

This document provides the following information about system configuration on the Extreme Networks 7100-Series platforms.

For information about...	Refer to page...
<a href="#">System Properties Overview</a>	9-1
<a href="#">User Management Overview</a>	9-5
<a href="#">Management Authentication Notification MIB Overview</a>	9-8
<a href="#">License Overview</a>	9-10
<a href="#">SNTP Overview</a>	9-11
<a href="#">Telnet Overview</a>	9-16
<a href="#">Secure Shell Overview</a>	9-17
<a href="#">Domain Name Server (DNS) Overview</a>	9-21
<a href="#">DHCP Overview</a>	9-23
<a href="#">DHCPv6 Overview</a>	9-28
<a href="#">Node Alias Overview</a>	9-34
<a href="#">MAC Address Settings Overview</a>	9-36
<a href="#">Terms and Definitions</a>	9-39

## System Properties Overview

Table 9-1 lists system parameter default values.

**Table 9-1 Default System Parameters**

Parameter	Description	Default Value
IP Gratuitous ARP	Provides an ARP announcement packet containing valid sender hardware and protocol addresses for the host that sent it.	disabled
System Utilization Threshold	Sets the threshold for sending CPU utilization notification messages.	800 (80%)
MTU	Sets the path MTU discovery protocol on the device.	enabled

You must configure your 7100-Series device with an IP interface and an IP address. You can also configure other system properties on the 7100-Series device. See [Table 9-2](#).

**Table 9-2 System Properties Configuration**

Task	Command
Set IP interfaces. You may specify an IP interface as the default management IP interface.	<b>set ip interface</b> <i>interface-name</i> [default]
Set the system IP address, subnet mask and default gateway. If not specified, <i>ip-mask</i> will be set to the natural mask of the <i>ip-address</i> and <i>ip-gateway</i> will be set to the <i>ip-address</i> . If not specified, the first IP interface configured on a system becomes the default IP interface.	<b>set ip address</b> <i>ip_address</i> [mask <i>ip_mask</i> ] [gateway <i>ip_gateway</i> ] [interface <i>interface-name</i> ]
Set the threshold for sending CPU utilization notification messages. A value of <b>0</b> will disable utilization notification messages.	<b>set system utilization threshold</b> <i>threshold</i>
Change the time of day on the system clock.	<b>set time</b> [ <i>mm/dd/yyyy</i> ] [ <i>hh:mm:ss</i> ]
Enable or disable the daylight savings time function.	<b>set summertime</b> {enable   disable} [ <i>zone</i> ]
Configure one of the following: <ul style="list-style-type: none"> <li>Specific dates to start and stop daylight savings time. These settings will be non-recurring and will have to be reset annually.</li> <li>Recurring daylight savings time settings. These settings will start and stop daylight savings time at the specified day of the month and hour each year and will not have to be reset annually.</li> </ul>	<b>set summertime date</b> <i>start_month start_date start_year start_hr_min end_month end_date end_year end_hr_min [offset_minutes]</i>  <b>set summertime recurring</b> <i>start_week start_day start_month start_hr_min end_week end_day end_month end_hr_min [offset_minutes]</i>
(Optional) Configure a name for the system. A name string containing a space in the text must be enclosed in quotes as shown in the example below. If <i>string</i> is not specified, the system name will be cleared.	<b>set system name</b> [ <i>string</i> ]
(Optional) Identify the location of the system. A location string containing a space in the text must be enclosed in quotes as shown in the example below. If <i>string</i> is not specified, the location name will be cleared.	<b>set system location</b> [ <i>string</i> ]
(Optional) Identify a contact person for the system. A contact string containing a space in the text must be enclosed in quotes as shown in the example below. If <i>string</i> is not specified, the contact name will be cleared.	<b>set system contact</b> [ <i>string</i> ]

**Table 9-2 System Properties Configuration (continued)**

Task	Command
Set the alias, a text name, for a physical object. If <i>string</i> is not specified, the specified alias will be cleared.	<b>set physical alias</b> {[chassis]   [module <i>module</i> ]   [powersupply <i>powersupply</i> ]   [powersupply-slot <i>powersupply-slot</i> ]   [poe-powersupply-poe-powersupply]   [fan <i>fantray</i> ]   [fan-slot <i>fantray</i> ]   [port <i>port-string</i> ]} [ <i>string</i> ]
Set the asset ID for a physical object.	<b>set physical assetid</b> {[chassis]   [module <i>module</i> ]   [powersupply <i>powersupply</i> ]   [poe-powersupply-poe-powersupply]   [fan <i>fantray</i> ]} <i>string</i>
Disable or re-enable path MTU discovery protocol on the device.	<b>set mtu</b> {enable   disable}

Table 9-3 lists system properties management and display commands for 7100-Series devices.

**Table 9-3 System Properties Management and Display Commands**

Task	Command
Display the gratuitous ARP processing behavior.	<b>show ip gratuitous-arp</b>
Display system information, including contact information, power and fan tray status and uptime.	<b>show system</b>
Display the system's hardware configuration.	<b>show system hardware</b>
Display system resource utilization information.	<b>show system utilization</b> [cpu   process   storage] [slot <i>slot</i> ]
Display the current time of day in the system clock.	<b>show time</b>
Display daylight savings time settings.	<b>show summertime</b>
Display the alias (a text name) for one or more physical objects.	<b>show physical alias</b> {[chassis]   [module <i>module</i> ]   [powersupply <i>powersupply</i> ]   [powersupply-slot <i>powersupply-slot</i> ]   [poe-powersupply-poe-powersupply]   [fan <i>fantray</i> ]   [fan-slot <i>fantray</i> ]   [port <i>port-string</i> ]}
Display the asset ID for a physical object.	<b>show physical assetid</b> {[chassis]   [powersupply <i>powersupply</i> ]   [poe-powersupply-poe-powersupply]   [fan <i>fan</i> ]}
Display the status of the path MTU (maximum transmission transmission) discovery protocol on the device.	<b>show mtu</b>
Display information about scheduled device resets.	<b>show reset</b>
Display output for technical support-related commands. Optionally, you can write this output to a file.	<b>show support</b> <i>filename</i>
Clear the IP interface.	<b>clear ip interface</b> <i>interface-name</i>
Clear an IP address.	<b>clear ip address</b> <i>ip-address</i>
Stop all gratuitous ARP processing.	<b>clear ip gratuitous-arp</b>
Clear the threshold for sending CPU utilization notification messages.	<b>clear system utilization</b>

**Table 9-3 System Properties Management and Display Commands (continued)**

Task	Command
Clear the daylight savings time configuration.	<b>clear summertime</b>
Reset the alias for a physical object to a zero-length string.	<b>clear physical alias</b> {[chassis]   [backplane backplane]   [slot slot]   [module module]   sub-module slot module   [powersupply powersupply]   [powersupply-slot powersupply-slot]   [poe-powersupply-poe-powersupply]   [fan fantray]   [fan-slot fantray]   [port port-string]}
Reset the asset ID for a module to a zero-length string.	<b>clear physical assetid</b> {[chassis]   [module module]   [powersupply powersupply]   [poe-powersupply-poe-powersupply]   [fan fan]}
Reset the state of the path MTU discovery protocol back to enabled.	<b>clear mtu</b>
Reset the device without losing any user-defined configuration settings or to display information about device resets.	<b>reset</b> {[mod   system] [cancel]}
Reset an option module CPU.	<b>reset nemcpu</b> mod.nemcpu
Schedule a system reset at a specific future time. This feature is useful for loading a new boot image.	<b>reset at</b> hh:mm [mm/dd] [reason]
Schedule a system reset after a specific time. This feature is useful for loading a new boot image.	<b>reset in</b> hh:mm [reason]
Clear all user-defined switch and router configuration parameters for one or all modules.	<b>clear config</b> mod_num   all

## System Properties Example

```

7100 Chassis(rw)->set ip interface vlan.0.5 default
7100 Chassis(rw)->set ip address 10.1.10.1 mask 255.255.128.0 gateway 10.1.10.1

7100 Chassis(rw)->set system utilization threshold 1000
7100 Chassis(rw)->set time 7:50:00
7100 Chassis(rw)->set summertime enable
7100 Chassis(rw)->set summertime recurring second Sunday March 02:00 first Sunday
November 02:00 60
7100 Chassis(rw)->set system name "Information Systems"
7100 Chassis(rw)->set system location "Bldg N32-04 Closet 9"
7100 Chassis(rw)->set system contact "Joe Smith"
7100 Chassis(rw)->set physical alias chassis chassison
7100 Chassis(rw)->set physical assetid module 1 bladel

```

## User Management Overview

An admin user (super user) can create user accounts, set the system password, and set the system lockout. Users with read-write access can change their own passwords. See [Procedure 9-1](#).

The 7100-Series device supports up to 16 user accounts, including the admin account, which cannot be disabled or deleted.

The 7100-Series supports security profiles that determine user access to certain commands and can also limit parameter settings for certain commands. The security profiles supported are normal and C2. The normal security profile provides standard user access based upon the configured user mode: super-user, read-write, and read-only. C2 is defined as Controlled Access Protection mode and is a security rating established by the U.S. National Computer Security Center (NCSC) and granted to products that pass Department of Defense (DoD) Trusted Computer System Evaluation Criteria (TCSEC) tests. A C2 rating ensures the minimum allowable levels of confidence demanded for government agencies and offices and other organizations that process classified or secure information. Use the **set security profile** command to set the security profile to either normal or C2 for the device. C2 mode can affect command availability and parameter value defaults and ranges. If C2 security mode affects a command, it is specified in the command entry found in the *Extreme Networks 7100-Series CLI Reference*.

Access to the boot menu during startup can be disabled. Access to the boot menu during startup is enabled by default.

The 7100-Series supports enabling of the Federal Information Processing Standards (FIPS) mode. FIPS mode is a mode where only FIPS approved authentication and encryption algorithms and methods are used. The current implementation supports the SHA1 algorithm in FIPS mode. Use the **set security fips mode** command to enable FIPS mode on the device.

User management configuration also includes the following:

- [“Setting the Authentication Login Method”](#) on page 9-7
- [“Using WebView”](#) on page 9-8

### Procedure 9-1 User Management Configuration

Step	Task	Command(s)
1.	Create a new user login account, or disable or enable an existing account.	<b>set system login</b> <i>username</i> [ <b>read-write</b>   <b>read-only</b>   <b>super-user</b> ] [ <b>enable</b>   <b>disable</b> ] [ <b>password</b> { <i>password</i>   <b>aging</b> { <i>days</i>   <b>disable</b>   <b>system</b> }}] [ <b>allowed-interval</b> { <i>HH:MM</i>   <i>HH:MM</i> }] [ <b>allowed-days</b> {[ <b>Sun</b> ] [ <b>Mon</b> ] [ <b>Tue</b> ] [ <b>Wed</b> ] [ <b>Thu</b> ] [ <b>Fri</b> ] [ <b>Sat</b> ]}] [ <b>simultaneous-logins</b> <i>num</i> ] [ <b>local-only</b> { <b>yes</b>   <b>no</b> }]
2.	Change system default passwords or set a new login password on the CLI. (Only available to users with super-user access.)	<b>set password</b> [ <i>username</i> ]

**Procedure 9-1 User Management Configuration (continued)**

Step	Task	Command(s)
3.	Configure system password parameters. A system password can contain the following special characters: !@#\$%^&*()-=[]\;?./`	<b>set system password</b> [aging {days   disable}] [history {size}] [length characters] [min-required-chars {[uppercase characters] [lowercase characters] [numeric characters] [special characters]}] [require-at-creation {yes   no}] [allow-duplicates {yes   no}] [allow-user-id {yes   no}] [substring-match-len characters] [allow-repeating-chars {num   yes   no}] [change-first-login {yes   no} [all]] [change-frequency minutes [all]] [expire-warning days] [grace-period {logins num   time days}]
4.	Optionally, disable access to the boot menu during bootup. Access to the boot menu is enabled by default.	<b>set security boot-access</b> {enable   disable}
5.	Set the number of failed login attempts before locking out (disabling) a read-write or read-only user account, the number of minutes to lockout the default admin super user account after maximum login attempts, and the number of inactive days before a non-superuser account is locked out.  If you set <b>inactive</b> to 0, no accounts will be locked out due to inactivity.  Once a user account is locked out, it can only be re-enabled by a super user with the <b>set system login</b> command.	<b>set system lockout</b> {[attempts attempts] [time minutes [all]] [port {enable   disable} [inactive days [all]] [emergency-access]}
6.	Optionally, enable FIPS mode on the device. Fips mode is disabled by default.	<b>set security fips mode</b> {enable   disable}
7.	Optionally, set the device's security profile. The security profile defaults to normal.	<b>set security profile</b> {c2   normal}

[Table 9-4](#) lists user account management and display commands for 7100-Series devices.

**Table 9-4 User Account Management and Display Commands**

Task	Command
To display user login account information.	<b>show system login [-verbose]</b>
To display current password configuration settings.	<b>show system password</b>
To display settings for locking out users.	<b>show system lockout</b>
To display the current boot access state for this device.	<b>show security boot-access</b>
To display the current security FIPS mode state for this device.	<b>show security fips mode</b>
To display the current security profile for this device.	<b>show security profile</b>

**Table 9-4 User Account Management and Display Commands (continued)**

Task	Command
To remove a local login user account or to reset a specified option to its default value. The account is removed if no optional parameters are entered.	<b>clear system login</b> <i>username</i> [allowed-interval] [allowed-days] [password [aging]] [simultaneous-logins] [local-only]
To reset system lockout parameters to default values.	<b>clear system lockout</b> [attempts] [time] [inactive]
To clear local login password parameters to default values. If no options are specified, all options are reset to default values.	<b>clear system password</b> [aging] [history] [length] [min-required-chars {[uppercase] [lowercase] [numeric] [special]}] [require-at-creation] [allow-duplicate] [allow-user-id] [substring-match-len] [allow-repeating-chars] [change-first-login] [change-frequency] [expire-warning] [grace-period]
To reset access to the boot menu during bootup to the default state of enabled.	<b>clear security boot-access</b>
To reset FIPS mode state to the default value of disabled on the device.	<b>clear security fips mode</b>
To reset the device security profile to the default value of normal.	<b>clear security profile</b>

## User Management Example

This example includes the following:

- Configuring system password parameters
- Creating a new user account
- Setting the password for the new user account
- Setting the system lockout parameters

```
7100 Chassis(su)->set system password age 60 length 6 allow-repeating-chars no
7100 Chassis(su)->set system login netops read-write enable
7100 Chassis(su)->set password rw
Please enter new password: *****
Please re-enter new password: *****
Password changed.
7100 Chassis(su)->set system lockout attempts 5 time 30 inactive 60
```

## Setting the Authentication Login Method

By default, the authentication login method is set to any, which uses the following precedence order:

- TACACS+
- RADIUS
- Local



**Procedure 9-2 Authentication Configuration**

Step	Task	Command(s)
1.	Change the default authentication login method. This command is not available to Read-Write users while in C2 security mode.	<b>set authentication login {any   local   radius   tacacs}</b>
2.	Display the current authentication login method to verify your changes.	<b>show authentication login</b>
3.	If necessary, reset the authentication login method to the default setting (any). This command is not available to Read-Write users while in C2 security mode.	<b>clear authentication login</b>
4.	Configure the chosen authentication login method. For more information, see <a href="#">Chapter 37, Security Configuration</a> for TACACS+ and <a href="#">Chapter 43, Authentication Configuration</a> for RADIUS.	

## Using WebView

By default, WebView (Extreme Networks' embedded web server for device configuration and management tasks) is enabled on TCP port number 80 of the 7100-Series device. You can verify WebView status, enable or disable WebView, and reset the WebView port.

[Procedure 9-3](#) describes how to configure WebView on an 7100-Series device.

**Procedure 9-3 WebView Configuration**

Step	Task	Command(s)
1.	Enable WebView	<b>set webview {enable   disable}</b>
2.	If necessary, change the TCP port for WebView from the default (port 80).	<b>set webview port <i>port</i></b>
3.	Display WebView status to verify your changes.	<b>show webview</b>

## Management Authentication Notification MIB Overview

You can enable or disable the sending of SNMP notifications when a user login authentication event occurs for various management access types. The types of access currently supported by the MIB include console, telnet, ssh, and web. By default, all Management Authentication Notification types are enabled.



**Note:** Ensure that SNMP is correctly configured in order to send these notifications. For more information, see [Chapter 17, Simple Network Management Protocol \(SNMP\) Configuration](#).

## Configuring Management Authentication Notification MIB

[Procedure 9-4](#) describes how to configure the Management Authentication Notification MIB on an 7100-Series device. Management Authentication Notification MIB commands can be entered in any command mode.

By default, all Management Authentication Notification types are enabled.

#### Procedure 9-4 Management Authentication Notification MIB Configuration

Step	Task	Command(s)
1.	Enable or disable the Management Authentication Notification MIB. By selecting the optional Management access type, you can specifically enable or disable a single access type, multiple access types or all of the access types.	<b>set mgmt-auth-notify {enable   disable} [console] [ssh] [telnet] [web]</b>
2.	Display the current setting for the Management Authentication Notification MIB.	<b>show mgmt-auth-notify</b>
3.	If necessary, set the current setting for the Management Authentication Notification access types to the default setting of enabled.	<b>clear mgmt-auth-notify</b>

## Management Authentication Notification MIB Configuration Examples

This example shows how to set all the authentication types to be disabled on the Management Authentication Notification MIB. That information is then displayed with the **show** command:

```
7100 Chassis(su)->set mgmt-auth-notify disable
7100 Chassis(su)->show mgmt-auth-notify
```

```
Management Type  Status
-----
console          disabled
ssh              disabled
telnet           disabled
web              disabled
```

This example shows how to set only the console and telnet authentication access types to be enabled on the Management Authentication Notification MIB. That information is then displayed with the **show** command:

```
7100 Chassis(su)->set mgmt-auth-notify enable console telnet
7100 Chassis(su)->show mgmt-auth-notify
```

```
Management Type  Status
-----
console          enabled
ssh              disabled
telnet           enabled
web              disabled
```

This example displays the state of Management Authentication Notification access types prior to using the **clear** command, then displays the same information after using the **clear** command:

```
7100 Chassis(su)->show mgmt-auth-notify
```

```
Management Type  Status
```

```

-----
console      enabled
ssh          disabled
telnet       enabled
web          disabled

7100 Chassis(su) ->clear mgmt-auth-notify

7100 Chassis(su) ->show mgmt-auth-notify

Management Type  Status
-----
console          enabled
ssh              enabled
telnet           enabled
web              enabled

```

## License Overview

A license, purchased separately, is available for the following:

- Advanced router for 7100-Series:
  - 71A-EOS-ADVL3 – 7100K-Series Advanced Routing license
  - 71A-EOS-G-ADVL3 – 7100G-Series Advanced Routing license

You must activate the purchased license key.

The license is activated on an 7100-Series module or chassis, as applicable, by using the **set license** command in any command mode to specify the license type and the ASCII advanced licensing key.

Use the **show license** command in any command mode to display the license key once you have activated the license.

## Configuring a License

[Procedure 9-5](#) describes how to configure the license on an 7100-Series device. License commands can be entered in any command mode.

### Procedure 9-5 License Configuration

Step	Task	Command(s)
1.	Activate the license on an 7100-Series device—module or chassis—as applicable. 7100-Series license keys can contain white spaces; therefore, you should enclose your license key in double quotation marks.	<b>set license advanced</b>
2.	Display the license key.	<b>show license</b>

## SNTP Overview

Simple Network Time Protocol (SNTP) provides for the synchronizing of system time for managed devices across a network. The 7100-Series implementation supports unicast polling and broadcast listening modes of operation to obtain the time from an SNTP server. SNTP is a subset of the Network Time Protocol (NTP) as specified in RFC 1305. The most recent version of SNTP is specified in RFC 2030. Since SNTP is a subset of NTP, all NTP servers are capable of servicing SNTP clients. The SNTP mode is set on the client using the **set sntp client** command.

### Unicast Polling Mode

When an SNTP client is operating in unicast mode, SNTP update requests are made directly to a server, configured using the **set sntp server** command. The client queries these configured SNTP servers at a fixed poll-interval configured using the **set sntp poll-interval** command. The order in which servers are queried is based on a precedence value optionally specified when you configure the server. The lower the configured precedence value, the higher the precedence for that server. The default is for all servers to have the same precedence. In this case, the server ordering is based upon the indexing of the server table.

The SNTP client makes a request to the SNTP server. The client waits a period of time configured using the **set sntp poll-timeout** command for a response from the server. If the poll timeout timer expires, the client will resend another request, up to the number of retries specified by the **set sntp poll-retry** command. If the retries have been exhausted, the client request is sent to the next server with the lowest configured precedence value or the next server in the server table, if precedence values are the same. If no server responds, the client waits the configured poll-interval time period and the process starts over again.

### Broadcast Listening Mode

With SNTP configured for broadcast listening mode, the client is passive and it is the broadcast server that broadcasts the time to the client. Broadcast listening uses the same poll-interval, poll-timeout and poll-retry values as unicast polling but they function differently. To account for the propagation delay between the server and the client, a broadcast delay value in milliseconds is configurable using the **set sntp broadcastdelay** command. The broadcast delay is the time window within which the device can accept a Broadcast SNTP packet from the SNTP server. Once the broadcast delay time window has ended, the poll interval window takes effect where the device will not accept Broadcast SNTP packets. When the poll interval window ends, the broadcast delay window starts again; SNTP packets can once again be accepted. If no Broadcast SNTP packets are seen within that broadcast delay window it is considered a timeout.

### SNTP Authentication

SNTP authentication provides the means for the SNTP client to authenticate the SNTP server using symmetric key cryptography. Because SNTP packet data is not sensitive information, the packet itself does not require encryption. Symmetric key cryptography uses a secret password shared between the SNTP client and server to generate an encrypted checksum which is appended to the SNTP packet data. The 7100-Series SNTP authentication supports 128-bit MD5 symmetric key cryptography.

SNTP authentication is configured by:

- Globally enabling the mode for the SNTP client
- Configuring up to 32 SNTP authentication key instances, by specifying:
  - A numeric key that identifies this SNTP authentication instance

- The MD5 authentication type
- A password as either an ASCII string of up to 32 printed characters (no white space) or the Hex formatted cypher produced by the previously entered ASCII string
- Associating an SNTP key instance with the SNTP server
- Enabling the authentication trust flag for the SNTP instance key assigned to the SNTP client

## Authentication Mode

SNTP authentication mode must be set to enabled for SNTP authentication to occur between the SNTP client and server. When the mode is set to enable, the SNTP client authenticates with the SNTP server before synchronization occurs. When the mode is set to disable, no authentication is performed on SNTP communications. SNTP authentication is set to disabled by default.

Use the **set sntp authentication mode** command to enable SNTP authentication on the SNTP client.

This example shows how to enable SNTP authentication mode:

```
7100 Chassis(rw)->set sntp authentication mode enable
```

## Authentication Key

The SNTP authentication key specifies the authentication instance to be used by the SNTP client when authenticating with the SNTP server. The SNTP client supports the configuration of up to 32 authentication keys. The authentication key instance ID is a numeric value. Each authentication key instance specifies the authentication type and password. SNTP authentication supports the MD5 authentication algorithm. The password is known to both the SNTP client and server. The password consists of an ASCII string of up to 32 non-white characters or the hexadecimal formatted cypher that was generated from the previously entered ASCII string.

Use the **set sntp authentication key** command to configure an authentication key instance.

This example shows how to create SNTP authentication key instances 1 - 3:

```
7100 Chassis(rw)->set sntp authentication key 1 md5 foobaraboof
7100 Chassis(rw)->set sntp authentication key 2 md5
DEADBFAFCFAFEBABEDEADBFAFCFAFEBAE
7100 Chassis(rw)->set sntp authentication key 3 md5
0123456789012345678901234567890
```

The SNTP authentication key is associated with an SNTP server using the **set sntp server** command.

This example shows how to set the server at IP address 10.21.1.100 as an SNTP server and to SNTP authenticate using authentication key instance 1:

```
7100 Chassis(rw)->set sntp server 10.21.1.100 key 1
```

## Authentication Trust Flag

The authentication trust flag specifies whether the key associated with it is enabled or disabled. When an authentication key trust flag is enabled, authentication will occur between the client and server the key is assigned to. If an authentication key trust flag is disabled, authentication will not occur between the client and server the key is assigned to.

The authentication trust flag is configured by specifying the instance the trust flag is associated with and whether the trust flag is enabled or disabled.

Use the **set sntp authentication trust** command to configure an SNTP authentication trust flag.

This example shows how to enable trust status for authentication key instance 1 and disable the trust status for authentication key instance 3:

```
7100 Chassis(rw)->set sntp authentication trust 1 enable
7100 Chassis(rw)->set sntp authentication trust 3 disable
```

## Configuring SNTP

This section provides details for the configuration of SNTP on the 7100-Series products.

[Table 9-5](#) lists SNTP parameters and their default values.

**Table 9-5 Default SNTP Parameters**

Parameter	Description	Default Value
SNTP authentication mode	Specifies whether authentication for all SNTP client communications is enabled or disabled.	disabled
SNTP authentication trust	Specified whether the trust state of an existing SNTP authentication key is enabled or disabled. Must be enabled for the SNTP authentication to occur.	disabled
SNTP mode	Specifies whether the current SNTP state is broadcast, unicast, or disabled.	disabled
unicast server precedence	Specifies a value that determines the order in which SNTP servers are polled if the precedence values are not the same.	1 (highest precedence)
broadcast delay	Specifies the propagation delay added to the time sent to the client in broadcast listening mode.	3000 milliseconds
poll-interval	Specifies the interval between unicast SNTP requests by the client to the server.	16 seconds
poll-retry	Specifies the number of times the client will resend the SNTP request to the server before moving on to the next server.	1
poll-timeout	Specifies the amount of time a client will wait for a response from the the SNTP server before retrying.	5 seconds
timezone offset	Specifies the offset in hours and minutes from UTC for this device	0 hours, 0 minutes

[Procedure 9-6](#) describes how configure SNTP. SNTP can be configured in any command mode.

### Procedure 9-6 Configuring SNTP

Step	Task	Command(s)
1.	Set the SNTP operation mode on the client.	<b>set sntp client {broadcast   unicast   disable}</b>

**Procedure 9-6 Configuring SNTP (continued)**

Step	Task	Command(s)
2.	When operating in broadcast mode, optionally change the broadcast delay period in milliseconds to be added to the server time for this client.	<b>set sntp broadcastdelay</b> <i>time</i>
3.	When operating in unicast mode, set the SNTP server(s) for this client, optionally specifying a precedence value per server.	<b>set sntp server</b> <i>ip-address</i> [ <i>precedence</i> ][ <b>key</b> <i>key-instance</i> ]
4.	When operating in unicast mode, optionally change the poll interval between SNTP unicast requests.	<b>set sntp poll-interval</b> <i>interval</i>
5.	When operating in unicast mode, optionally change the number of poll retries to a unicast SNTP server.	<b>set sntp poll-retry</b> <i>retry</i>
6.	When operating in unicast mode, optionally change the poll timeout for a response to a unicast SNTP request.	<b>set sntp poll-timeout</b> <i>timeout</i>
7.	Optionally, set the SNTP time zone name and the hours and minutes it is offset from Coordinated Universal Time (UTC).  <b>Note:</b> The daylight savings time function can be enabled and associated with the timezone set here using the <b>set summertime</b> command.	<b>set timezone</b> <i>name</i> [ <i>hours</i> ] [ <i>minutes</i> ]
8.	Optionally, enable authentication for all SNTP client communications.	<b>set sntp authentication mode</b> { <b>enable</b>   <b>disable</b> }
9.	Optionally, create a new or modify an existing SNTP authentication key.	<b>set sntp authentication key</b> <i>key-instance</i> <i>type</i> <i>password</i>
10.	Optionally, change the SNTP authentication trust state for an authentication key.	<b>set sntp authentication trust</b> <i>key-instance</i> { <b>enable</b>   <b>disable</b> }

[Table 9-6](#) describes how to manage and display SNTP.

**Table 9-6 Managing and Displaying SNTP**

Task	Command(s)
To display SNTP client settings:	<b>show sntp</b>
To set the SNTP client's operational mode to disable:	<b>clear sntp client</b>
To remove one or all servers from the SNTP server list:	<b>clear sntp server</b> { <i>ip-address</i>   <b>all</b> }
To reset the delay time for SNTP broadcast frames to its default value:	<b>clear sntp broadcastdelay</b>
To reset the poll interval between unicast SNTP requests to its default value:	<b>clear sntp poll-interval</b>
To reset the number of poll retries to a unicast SNTP server to its default value:	<b>clear sntp poll-retry</b>
To reset the SNTP poll timeout to its default value:	<b>clear sntp poll-timeout</b>
To display the current timezone setting:	<b>show timezone</b>
To remove the SNTP timezone adjustment values:	<b>clear timezone</b>

**Table 9-6 Managing and Displaying SNTP (continued)**

Task	Command(s)
To clear SNTP authentication key configuration or reset the SNTP authentication mode to the default value:	<b>clear sntp authentication {all   key key-instance   mode}</b>

## SNTP Configuration Examples

The following example configures the client for SNTP broadcast mode:

- Setting the broadcast delay to 3500 milliseconds
- Setting the timezone to Eastern Daylight Time (EDT)
- Displaying the current SNTP configuration

```
7100 Chassis(rw)->set sntp client broadcast
7100 Chassis(rw)->set sntp broadcastdelay 3500
7100 Chassis(rw)->set timezone EDT -4 0
7100 Chassis(rw)->show sntp
```

```
SNTP Version: 4
Current Time: SAT AUG 01 14:34:53 2009
Timezone: 'EDT', offset from UTC is -4 hours and 0 minutes
Client Mode: broadcast
Broadcast Delay: 3500 microseconds
Broadcast Count: 1
Poll Interval: 512 seconds
Poll Retry: 1
Poll Timeout: 5 seconds
SNTP Poll Requests: 0
Last SNTP Update: SAT AUG 01 14:23:54 2009
Last SNTP Request: SAT AUG 01 14:23:54 2009
Last SNTP Status: Enabled
```

```
Status          Precedence      SNTP-Server
-----
Active          1               10.21.1.300
7100 Chassis(rw)->
```

The following example configures the client for SNTP unicast mode with SNTP authentication operational:

- Enables SNTP authentication mode
- Creates an SNTP authentication key instance 1 and sets the password to foobar
- Sets the SNTP server to IP address 10.21.1.100 and assigns authentication key instance 1 to it
- Set the SNTP authentication key trust flag to enable for key instance 1
- Sets the SNTP poll interval to 600 seconds
- Sets the UTC timezone to Eastern Daylight Time (EDT)



- Sets the poll retry to 2
- Displays the current SNTP configuration

```

7100 Chassis(rw)->set sntp client unicast
7100 Chassis(rw)->set sntp authentication mode enable
7100 Chassis(rw)->set sntp authentication key 1 md5 foobar
7100 Chassis(rw)->set sntp authentication trust 1 enable
7100 Chassis(rw)->set sntp server 10.21.1.100 key 1
7100 Chassis(rw)->set sntp poll-interval 600
7100 Chassis(rw)->set timezone EDT -4 0
7100 Chassis(rw)->set sntp poll-retry 2
7100 Chassis(rw)->show sntp
SNTP Version: 4
Current Time: FRI MAY 06 15:33:53 2011
Timezone: 'EDT', offset from UTC is -4 hours and 0 minutes
Client Mode: unicast
Broadcast Delay: 3000 microseconds
Broadcast Count: 0
Poll Interval: 600 seconds
Poll Retry: 2
Poll Timeout: 5 seconds
SNTP Poll Requests: 2
Last SNTP Update: MON MAY 02 14:42:52 2011
Last SNTP Request: MON MAY 02 14:42:52 2011
Last SNTP Status: Enabled

```

SNTP Servers:

Status	Precedence	Key	SNTP-Server
Active	1	1	10.21.1.100

SNTP Authentication: Enabled

Status	Key	Type	Trusted
Active	1	MD5	Enabled

```
7100 Chassis(rw)->
```

## Telnet Overview

Telnet provides an unsecured communications method between a client and the switch.

Telnet is activated by enabling Telnet on the device, using the **set telnet enable** command in any command mode.

Use the **show telnet** command in any command mode to display whether Telnet is currently enabled or disabled.

## Configuring Telnet

[Procedure 9-7](#) describes how to configure and use Telnet on an 7100-Series device. Telnet commands can be entered in any command mode.

### Procedure 9-7 Telnet Configuration

Step	Task	Command(s)
1.	Enable or disable either inbound or outbound or both Telnet services.	<b>set telnet {enable   disable} {all   inbound   outbound}</b>
2.	Verify the Telnet status.	<b>show telnet</b>
3.	Start a Telnet connection. <ul style="list-style-type: none"> <li>• <b>-s</b> - The source IP address to use in the outgoing telnet</li> <li>• <b>-4   -6</b> - Use only IPv4 or IPv6 addresses but not both</li> <li>• <b>-vrf</b> - The name of the router used for this session</li> <li>• <b>-r</b> - Bypass the host routing table for this session</li> <li>• <i>host</i> - The remote host to Telnet to for this session</li> </ul>	<b>telnet [-s <i>src-addr</i>] [-4   -6] [-vrf <i>router</i>] [-r] {<i>host</i> [<i>port</i>]}</b>

## Telnet Examples

The following example shows how to enable Telnet:

```
7100 Chassis(rw)->set telnet enable all
```

The following example shows how to verify the Telnet status:

```
7100 Chassis(rw)->show telnet
```

```
Telnet inbound is currently: ENABLED
```

```
Telnet outbound is currently: ENABLED
```

The following example telnets to remote host 10.21.42.01:

```
7100 Chassis(rw)->telnet 10.21.42.01
```

## Secure Shell Overview

The Secure Shell (SSH) security feature provides a secure encrypted communications method between a client and the switch providing data privacy and integrity that is an alternative to the unsecure Telnet protocol. Using SSH, the entire session is encrypted, including the transmission of user names and passwords, and negotiated between a client and server both configured with the SSH protocol. Telnet sessions are unsecure. All data is sent unencrypted. Use SSH instead of Telnet when the security of login and data transmission is a concern.

The 7100-Series SSHv2 implementation includes:

- Data privacy
- Communication integrity

An SSH server resides on the 7100-Series platform and listens for client connection requests. Once a request is authenticated, a secure connection is formed through which all subsequent traffic is sent. All traffic is encrypted across the secure channel, which ensures data integrity. This prevents someone from seeing clear text passwords or file content, as is possible with the Telnet application.

Once SSH has been enabled and the 7100-Series has at least one valid IP address, you can establish an SSH client session from any TCP/IP based node on the network, by using an application supporting SSH to connect to an IP address and entering your user name and password. Refer to the instructions included with your SSH application for information about establishing a session.

SSH is activated by enabling the SSH server on the device, using the **set ssh enable** command in any command mode.

Enabling the server automatically generates a host key for the server, used during the life of the client to server connection. The host key type can be set to either dsa or rsa. The host key type defaults to rsa.

There is one host key per device; every time an SSH client logs into a device it should see the same host key; if the host key is different, the SSH Client warns you that the host key has changed. The following is a sample warning when an SSH Client detects a new host key:

```
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@   WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!   @
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@

IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!

Someone could be eavesdropping on you right now (man-in-the-middle attack)!

It is also possible that the RSA host key has just been changed.

The fingerprint for the RSA key sent by the remote host is
67:c6:71:ff:e8:02:7c:ce:0f:0d:67:67:63:a8:2e:9c.

Please contact your system administrator.

Add correct host key in /home/documentation/doc1/.ssh/known_hosts to get rid of
this message.

Offending key in /home/documentation/doc1/.ssh/known_hosts:24
RSA host key for 10.4.99.4 has changed and you have requested strict checking.
Host key verification failed.
```

The SSH server can be reinitialized. Reinitializing the server clears all current client to server connections. Reinitializing the server does not reinitialize the host key. Should you believe the host key has been compromised, or otherwise wish to change it, the host key can be reinitialized using the **set ssh hostkey reinitialize** command.

An SSH session to a remote host can be started using the **ssh** command.

Use the **show ssh state** command in any command mode to display whether SSH is currently enabled or disabled.

## SSH Client Authentication

The password authentication method is supported for a remote SSH client attempting to login to the 7100-Series SSH server.

## Password Authentication

Once SSH has been enabled and the 7100-Series has at least one valid IP address, you can establish an SSH client session from any TCP/IP based node on the network, by using an application supporting SSH to connect to an IP address and entering your user name and password. Refer to the instructions included with your SSH application for information about establishing a session.

## Configuring Secure Shell

[Procedure 9-8](#) describes how to configure Secure Shell on an 7100-Series device. Secure Shell commands can be entered in any command mode.

### Procedure 9-8 SSH Configuration

Step	Task	Command(s)
1.	Enable, disable, or reinitialize the SSH server.	<b>set ssh {enable   disable   reinitialize}</b>
2.	Optionally modify the SSH client alive interval.	<b>set ssh client alive-interval <i>interval</i></b>
3.	Optionally modify the the maximum number of times a client alive message will be sent before the session times out.	<b>set ssh client alive-count <i>count</i></b>
4.	Set or reinitialize the host key on the SSH server.	<b>set ssh hostkey [reinitialize] [type <i>type</i>]</b>

**Procedure 9-8 SSH Configuration (continued)**

Step	Task	Command(s)
5.	<p>Start an SSH session.</p> <p><i>hostname</i> - Specifies the host name or IP address of the remote host this SSH session is connecting to.</p> <p><b>-4   -6</b> - Optionally specifies that SSH should use either IPv4 or IPv6 addresses, but not both.</p> <p><b>-b bind-address</b> - Optionally specifies the IP address to transmit from when there are multiple interfaces and or addresses.</p> <p><b>-c cipher-spec</b> - Optionally specifies a list of the cipher specifications allowed for encrypting this session.</p> <p><b>-e escape-char</b> - Optionally sets the escape character for the session.</p> <p><b>-l login-name</b> - Optionally specifies the user to login as on the remote host.</p> <p><b>-m mac-spec</b> - Optionally specifies the MAC algorithms used for data integrity protection.</p> <p><b>-p port</b> - Optionally specifies the host port to connect to on the remote host.</p> <p><b>-q</b> - Optionally specifies that the session will operate in quiet mode, causing all warning and diagnostic messages to be suppressed.</p> <p><b>-r</b> - Optionally specifies that normal routing table lookup should be bypassed and that the session request should be sent directly to a host on an attached network.</p> <p><b>-v</b> - Optionally specifies that the session will operate in verbose mode, causing SSH to print debugging messages about its progress.</p> <p><b>-vrf router</b> - Optionally specifies the router on which to source this SSH session.</p>	<pre>ssh hostname [-4   -6] [-b bind-address] [-c cipher-spec] [-e escape-char] [-l login-name] [-m mac-spec] [-p port] [-q] [-r] [-v] [-vrf router]</pre>
6.	Verify the SSH state.	<b>show ssh state</b>

## Secure Shell Configuration Examples

The following commands enable and verify SSH:

```
7100 Chassis(rw)->set ssh enable
7100 Chassis(rw)->show ssh state
SSH Server state: Enabled
7100 Chassis(rw)->
```

The following command reinitializes the host key on the SSH server:

```
7100 Chassis(rw)->set ssh hostkey reinitialize
```

## Domain Name Server (DNS) Overview

The Domain Name Server (DNS) resolver is a session layer protocol that maps network host names to IP addresses (and vice versa). The client function queries configured servers to provide mapping services for CLI commands (for example, ping, telnet) which allow a hostname to be specified.

The DNS resolver feature is enabled by default. Up to four DNS servers can be configured for DNS resolution. The domain name (Net, Host, Gateway, or Domain name) associated with this device can be configured. A default DNS zone can be specified indicating the initial zone used for DNS lookup. Supported zones are IPv4 and IPv6. The default zone is IPv4. The default zone names are:

- IPv4: - **in-addr.arpa**
- IPv6: - **ip6.int**

The port number the DNS resolver uses for DNS queries can be configured. The default port is 53. DNS requests will time out and retry the request after a configurable number of seconds. After a configurable amount of retries, if there is more than a single DNS server configure, the request will be sent to the next configured server for up to the number of configured retries.

## Configuring DNS

This section provides details for the configuration of DNS resolution on the 7100-Series products.

[Table 9-7](#) lists DNS parameters and their default values.

**Table 9-7 Default DNS Parameters**

Parameter	Description	Default Value
DNS resolver state	Specifies whether DNS resolver is enabled or disabled on the device.	enabled
DNS zone	Specifies the DNS zone for IPv4 and IPv6.	IPv4 - in-addr.arpa IPv6 - ip6.arpa
DNS port	Specifies the port number the DNS resolver uses for DNS queries.	53
timeout	Specifies the number of seconds before a DNS request is retried when the DNS server fails to respond.	10 seconds
query-retries	Specifies the number of times to retry a lookup request to a DNS server that has failed to respond.	2

[Procedure 9-9](#) describes how to configure DNS resolution. DNS can be configured in any CLI command mode.

**Procedure 9-9 Configuring DNS Resolution**

Step	Task	Command(s)
1.	Enable DNS on the switch if you have manually disabled it. DNS is enabled by default.	<b>set ip dns enable</b>
2.	Optionally, set the domain name for this device.	<b>set ip dns domain name</b>
3.	Configure the DNS servers for this device. Valid server values are: <b>primary, secondary, tertiary, quaternary.</b>	<b>set ip dns server ip-address server</b>

**Procedure 9-9 Configuring DNS Resolution (continued)**

Step	Task	Command(s)
4.	Optionally, configure the DNS zone for IPv4 and IPv6 IP address to name lookups.	<b>set ip dns zone {ipv4   ipv6} zone-name</b>
5.	Optionally, configure the port number the DNS resolver uses for DNS queries. The default port is <b>53</b> .	<b>set ip dns port-number port-number</b>
6.	Optionally, change the number of seconds before a DNS request is retried when the DNS server fails to respond.	<b>set ip dns timeout seconds</b>
7.	Optionally, change the number of times to retry a lookup request to a DNS server that has failed to respond.	<b>set ip dns query-retries retries</b>

[Table 9-8](#) describes how manage DNS resolution on an 7100-Series switch. DNS commands can be configured in any CLI command mode.

**Table 9-8 Managing DNS Resolution**

Task	Command(s)
To clear the DNS domain name configuration.	<b>clear ip dns domain</b>
To clear the DNS server configuration.	<b>clear ip dns server [server   all]</b>
To reset the DNS IPv4 or IPv6 zone configuration.	<b>clear ip dns zone [ipv4   ipv6]</b>
To reset the DNS port number used for DNS queries to the default value.	<b>clear ip dns port-number</b>
To reset the DNS timeout to the default value.	<b>clear ip dns timeout</b>
To reset the number DNS query retries to the default value.	<b>clear ip dns query-retries</b>
To clear all DNS configuration to the default state.	<b>clear ip dns all</b>
To reset DNS status for this device to the default value.	<b>clear dns status</b>
To display DNS configuration for this device.	<b>show ip dns</b>

## DNS Configuration Example

The following DNS configuration example:

- Sets the DNS domain name to **Extremenetworks.Documentation**
- Configures two DNS servers:
  - Primary - **123.50.50.10**
  - Secondary - **123.50.50.20**
- Configures the DNS timeout value to **4** seconds
- Configures the number of query retries to **3**

```
7100-Series(rw)->set ip dns domain Extremenetworks.Documentation
7100-Series(rw)->set ip dns server 153.50.50.10 primary
7100-Series(rw)->set ip dns server 153.50.50.20 secondary
7100-Series(rw)->set ip dns timeout 4
```

```

7100-Series(rw)->set ip dns query-retries 3
7100-Series(rw)->show ip dns
Current State:                Enabled
Default DNS domain name:     Extremenetworks.Documentation
DNS zones:
  IPv4:                       in-addr.arpa
  IPv6:                       ip6.int

DNS port number:            53
DNS server timeout:        4 seconds
DNS query retries:        3
DNS Name servers                Status
-----
153.50.50.10                  primary
153.50.50.20                  secondary
7100-Series(rw)->

```

## DHCP Overview

The Dynamic Host Configuration Protocol (DHCP) provides services for allocating and delivering IPv4 addresses and other IPv4 DHCP server options to Internet hosts. DHCP consists of two components: a protocol for delivering host-specific configuration parameters from a DHCP server to a host, and a mechanism for allocating network addresses to hosts. Optional functionality also provides services to complete high-availability, authenticated and QoS-dependant host configuration.

The DHCP protocol is based on a client-server model in which a designated DHCP server allocates network addresses and delivers configuration parameters to dynamically configured clients. Throughout the remainder of this section, the term “server” refers to a host providing initialization parameters through DHCP, and the term “client” refers to a host requesting initialization parameters from a DHCP server.

DHCP supports the following mechanisms for IP address allocation:

- Automatic — DHCP assigns an IP address to a client for a limited period of time (or until the client explicitly relinquishes the address).
- Manual — A client's IP address is assigned by the network administrator, and DHCP is used simply to convey the assigned address to the client.

The amount of time that a particular IP address is valid for a system is called a lease. The 7100-Series device maintains a lease database which contains information about each assigned IP address, the MAC address to which it is assigned, the lease expiration, and whether the address assignment is dynamic or static. The DHCP lease database is stored in flash memory.



**Note:** The 7100-Series DHCP server is not designed to work as the primary DHCP server in an enterprise environment with hundreds of clients that are constantly seeking IP address assignment or reassignment. A standalone DHCP server with a redundant backup server may be more suitable for this type of environment.



## IPv4 DHCP Supported Server Options

Table 9-9 on page 9-24 lists the IPv4 DHCP server option names and codes supported by the firmware. All options specified in Table 9-9 may be configured using the **option** command. Several commonly-used options may also be configured using dedicated commands:

**domain-name**, **dns-server**, **netbios-name-server**, **netbios-node-type**, and **default-router**. These commands are specified in Procedure 9-11 on page 9-32

Except where noted, all options are defined in RFC-2132. In addition, the site-specific option codes designated by RFC-2132 (128-254) may be used to define options for use within a site or an organization. Some vendors have made use of site-specific options to configure their product features.

**Table 9-9 IPv4 DHCP Server Codes**

Description	Code	Input Methods	RFC
SubnetMask:	1	hex,ip	RFC 2132
TimeOffset:	2	hex	RFC 2132
Router:	3	hex,ip	RFC 2132
TimeServer:	4	hex,ip	RFC 2132
NameServer:	5	hex,ip	RFC 2132
DomainNameServer:	6	hex,ip	RFC 2132
LogServer:	7	hex,ip	RFC 2132
CookieServer:	8	hex,ip	RFC 2132
LPRServer:	9	hex,ip	RFC 2132
ImpressServer:	10	hex,ip	RFC 2132
ResourceLocationServer:	11	hex,ip	RFC 2132
HostName:	12	ascii,hex	RFC 2132
BootFileSize:	13	hex	RFC 2132
MeritDumpName:	14	ascii,hex	RFC 2132
DomainName:	15	ascii,hex	RFC 2132
SwapServer:	16	hex,ip	RFC 2132
RootPath:	17	ascii,hex	RFC 2132
ExtensionsPath:	18	ascii,hex	RFC 2132
IpForwarding:	19	hex	RFC 2132
NonLocalSourceRouting:	20	hex	RFC 2132
PolicyFilter:	21	hex,ip	RFC 2132
MaxDatagramReassemblySize:	22	hex	RFC 2132
DefaultIpTTL:	23	hex	RFC 2132
PathMTUAgingTimeout:	24	hex	RFC 2132
PathMTUPlateauTable:	25	hex	RFC 2132
InterfaceMTU:	26	hex	RFC 2132
AllSubnetsLocal:	27	hex	RFC 2132

**Table 9-9 IPv4 DHCP Server Codes**

Description	Code	Input Methods	RFC
BroadcastAddress:	28	hex,ip	RFC 2132
PerformMaskDiscovery:	29	hex	RFC 2132
MaskSupplier:	30	hex	RFC 2132
PerformRouterDiscovery:	31	hex	RFC 2132
RouterSolicitationAddress:	32	hex,ip	RFC 2132
StaticRoute:	33	hex,ip	RFC 2132
TrailerEncapsulation:	34	hex	RFC 2132
ARPCacheTimeout:	35	hex	RFC 2132
EthernetEncapsulation:	36	hex	RFC 2132
TCPDefaultTTL:	37	hex	RFC 2132
TCPDefaultKeepaliveInterval:	38	hex	RFC 2132
TCPDefaultKeepaliveGarbage:	39	hex	RFC 2132
NISDomain:	40	ascii,hex	RFC 2132
NISServers:	41	hex,ip	RFC 2132
NTPServers:	42	hex,ip	RFC 2132
VendorSpecificInfo:	43	ascii,hex	RFC 1533 RFC 2132
NetBIOSNameServer:	44	hex,ip	RFC 1533 RFC 2132
NetBIOSDatagramDistributionServer:	45	hex,ip	RFC 1533 RFC 2132
NetBIOSNodeType:	46	hex	RFC 1533 RFC 2132
NetBIOSScope:	47	ascii,hex	RFC 1533 RFC 2132
XWindowFontServer:	48	hex,ip	RFC 1533 RFC 2132
XWindowDisplayManager:	49	hex,ip	RFC 1533 RFC 2132
IpAddressLeaseTime:	51	hex	RFC 1533 RFC 2132
RenewalTimeValue:	58	hex	RFC 1533 RFC 2132
RebindingTimeValue:	59	hex	RFC 1533 RFC 2132
NISPlusDomain:	64	ascii,hex	RFC 2132
NISPlusServers:	65	hex,ip	RFC 2132

**Table 9-9 IPv4 DHCP Server Codes**

Description	Code	Input Methods	RFC
TFTPServerName:	66	ascii,hex	RFC 2132
BootfileName:	67	ascii,hex	RFC 2132
MobileIpHomeAgent:	68	hex,ip	RFC 2132
SMTPServer:	69	hex,ip	RFC 2132
POP3Server:	70	hex,ip	RFC 2132
NNTPServer:	71	hex,ip	RFC 2132
DefaultWWWServer:	72	hex,ip	RFC 2132
DefaultFingerServer:	73	hex,ip	RFC 2132
DefaultIRCServer:	74	hex,ip	RFC 2132
StreetTalkServer:	75	hex,ip	RFC 2132
STDAserver:	76	hex,ip	RFC 2132
UserClass:	77	ascii,hex	RFC 3004
SLPDirectoryAgents:	78	hex	RFC 2610
SLPServiceScope:	79	hex	RFC 2610

## DHCP Server

DHCP provides the following mechanisms for IP address allocation by a DHCP server:

- **Automatic**—DHCP assigns an IP address, from a range of addresses defined by the **ip local pool** command in configuration mode and configured as a pool of addresses by the **ip dhcp pool** command. The address is assigned to a client for a limited period of time set by the **lease** command (or until the client explicitly relinquishes the address). The **exclude** command is used to exclude one or more IP addresses from a DHCP local address pool.
- **Manual**—A client's IP address is assigned by the network administrator using the **host** command in DHCP host configuration command mode, and DHCP is used simply to convey the assigned address to the client. Enter DHCP host configuration command mode using the **hardware-address** or **client-identifier** commands in DHCP pool configuration command mode. The **hardware-address** or **client-identifier** command specifies the client hardware address and client unique identifier, respectively.

The 7100-Series device maintains a lease database which contains information about each assigned IP address, the MAC address/unique identifier to which it is assigned, the lease expiration, and whether the address assignment is automatic or static.

In addition to assigning IP addresses, the DHCP server can also be configured to assign the following to requesting clients:

- Default router(s), using the **default-router** command in DHCP pool configuration command mode
- DNS server(s), using the **dns-server** command, and domain name, using the **domain-name** command in DHCP pool configuration command mode
- NetBIOS WINS server(s), using the **netbios-name-server** command, and node type, using the **netbios-node-type** command in DHCP pool configuration command mode

- Boot file, using the **bootfile** command mode in DHCP pool configuration command mode
- DHCP options as defined by RFC 2132, using the **option** command in DHCP pool configuration command mode
- Next server in the DHCP server boot process, using **next-server** in the DHCP pool configuration command mode

## Configuring Client Class

DHCP client class provides a logical container for a set of client properties, allowing the assignment of a client property set to a DHCP client rather than configuring each client separately. Client-classes are created within a DHCP pool context using the **client-class** command. There are two modes in which a client-class can be assigned, by:

- Directly associating a client-class with a client binding using either the **hardware-address** or **client-identifier** commands
- Receiving a dynamic request with DHCP option 77 (user class) client-class match

## DHCP Configuration Example

In the following example client-class **class1** will be configured with a default router 3.3.3 and a DNS server 4.4.4.4. When we assign client-class **class1** to client 00:11:22:33:44:55 using the **hardware-address** command within DHCP pool **pool1**, the pool settings for default router (1.1.1.1) will be overwritten by the client-class **class1** settings for this client and any client that should receive a dynamic request with DHCP option 77 specifying client-class **class1**. The DNS server setting will be neither the **pool1** setting nor the **class1** setting. It will be manually set in the host configuration mode for this client to IP address 5.5.5.5. If it were not manually set, it would take the setting specified in **class1**.

**pool1** settings also include:

- Domain name of MyCompany.com
- Boot file: dhcpboot
- The assigning of WWW servers 10.70.0.10 10.70.0.11 10.70.0.12 to this pool using option 72 (WWW servers)
- A DHCP boot process next server: 10.70.0.12
- A pool lease of 100 days

These settings will apply to any client configured within pool1 that is not overwritten by either a client class setting or a received option setting.

The example first configures a local pool **pool1** to either automatically or allow the manual setting of IP addresses from the 10.60.0.0 subnet. IP addresses 10.60.0.10 - 30 are excluded from the local **pool1**. These addresses cannot be automatically or manually assigned to clients in this pool. DHCP pool configuration is then entered for **pool1** setting the default router to 1.1.1.1 and the DNS server to 2.2.2.2. When client classes are not applied, these values will be configured along with all the other values listed for this pool.

Client-class class1 is configured as specified above. The client-class class1 is applied to client 00:11:22:33:44:55. Entering host configuration mode for this client, the DNS server is set to IP address 5.5.5.5. This setting will override the class1 DNS server setting for this client. The host IP address for this client is manually set to 10.60.0.1 from the local pool. If the client IP address were not manually set, the client IP address would have been automatically set from the local pool of addresses configured for **pool1**.

```
7100 Chassis(rw-config)->ip local pool pool1 10.60.1.0 255.255.255.0
7100 Chassis(rw-config-ip-local-pool)->exclude 10.60.1.10 20
```

```
7100 Chassis(rw-config-ip-local-pool)->exit
7100 Chassis(rw-config)->ip dhcp pool pool1
7100 Chassis(rw-config-dhcp-pool)->domain-name MyCompany.com
7100 Chassis(rw-config-dhcp-pool)->bootfile dhcpboot
7100 Chassis(rw-config-dhcp-pool)->option 72 ip 10.70.0.10 10.70.0.11 10.70.0.12
7100 Chassis(rw-config-dhcp-pool)->next-server 10.70.0.12
7100 Chassis(rw-config-dhcp-pool)->lease 100
7100 Chassis(rw-config-dhcp-pool)->default-router 1.1.1.1
7100 Chassis(rw-config-dhcp-pool)->dns-server 2.2.2.2
7100 Chassis(rw-config-dhcp-pool)->client-class class1
7100 Chassis(rw-config-dhcp-class)->default-router 3.3.3.3
7100 Chassis(rw-config-dhcp-class)->dns-server 4.4.4.4
7100 Chassis(rw-config-dhcp-class)->exit
7100 Chassis(rw-config-dhcp-pool)->hardware-address 00:11:22:33:44:55
client-class      class1
7100 Chassis(rw-config-dhcp-host)->dns-server 5.5.5.5
7100 Chassis(rw-config-dhcp-host)->host 10.60.0.1
7100 Chassis(rw-config-dhcp-host)->exit
7100 Chassis(rw-config-dhcp-pool)->exit
7100 Chassis(rw-config)->
```

## DHCPv6 Overview

The IPv6 Dynamic Host Configuration Protocol (DHCPv6) provides services for delivering DHCPv6 server options to requesting clients. DHCPv6 options are contained in a DHCPv6 pool. The pool is assigned to the DHCPv6 server to which the client requests server options information.

Use the **ipv6 dhcp pool** command in global router configuration mode to create the DHCPv6 server pool and enter pool configuration mode.

Use the **ipv6 dhcp server** command in interface configuration mode to assign the pool to the DHCPv6 server. The assigned pool must already be created before you assign the pool to a DHCPv6 server.

Use this **domain-name** command to return one or more domain names when responding to a DHCPv6 client request.

Use the **dns-server** command to assign one or more DNS servers to DHCPv6 clients.

Use the **nis-domain-name** command to return one or more Network Information Services (NIS) domain names when responding to a DHCPv6 client request.

Use the **nis-server** command to assign one or more Network Information Services (NIS) servers to DHCPv6 clients.

Use the **nisp-domain-name** command to return one or more Network Information Services (NIS) version 2 domain names when responding to a DHCPv6 client request.

Use the **nisp-server** command to assign one or more Network Information Services (NIS) version 2 servers to DHCPv6 clients.

Use the **sip-domain-name** command to return one or more Session Initiation Protocol (SIP) domain names when responding to a DHCPv6 client request.

Use the **sip-server** command to assign one or more Session Initiation Protocol (SIP) servers to DHCPv6 clients.

Use the **sntp-server** command to assign a Simple Network Time Protocol (SNTP) server to DHCPv6 clients.

Use the **unicast-server** command to assign a unicast server to DHCPv6 clients.

Use the **information-refresh** command to configure the amount of time a client should wait before refreshing information from the DHCPv6 server.

**Table 9-10 DHCPv6 Server Supported Options**

DHCPv6 Option	Option Code	RFC
Unicast Server	12	RFC 3315
SIP Domain Name	21	RFC 3319
SIP Server	22	RFC 3319
DNS Server	23	RFC 3646
Domain Search List	24	RFC 3646
NIS Server	27	RFC 3898
NISP server	28	RFC 3898
NIS Domain Name	29	RFC 3898
NISP Domain Name	30	RFC 3898
SNTP Server	31	RFC 4075
Information Refresh Time	32	RFC 4242

## DHCPv6 Server Option Information Configuration Example

This example enables the DHCPv6 server on VLAN 1 with the DHCPv6 option information pool docPool by:

- Creating the docPool DHCPv6 option information pool and entering DHCPv6 option configuration mode and assigning the following options to the pool:
  - The **myEnterprise.com** domain name
  - An IPv6 DNS server at addresses **1111::12**, **1111::13**, and **1111::14**
  - The **myNisEnterprise.com** NIS domain name
  - An IPv6 NIS-DNS server at addresses **1111::12**, **1111::13**, and **1111::14**
  - The **myNispEnterprise.com** NISP domain name
  - An IPv6 NISP-DNS server at addresses **1111::12**, **1111::13**, and **1111::14**
  - The **mySipEnterprise.com** SIP domain name
  - An IPv6 SIP-DNS server at addresses **1111::12**, **1111::13**, and **1111::14**
  - An SNTP server at address **1111::15**
  - A unicast server at address **1111::15**
  - A client information refresh wait of **12** hours
- Configuring the DHCP server on VLAN **44** with a **dhcpPool1** DHCPv6 pool

```
7100 Chassis(rw) ->configure
7100 Chassis(rw-config) ->ipv6 dhcp pool docPool
7100 Chassis(rw-config-dhcp-v6-pool) ->domain-name myEnterprise.com
7100 Chassis(rw-config-dhcp-v6-pool) ->dns-server 1111::12 1111::13 1111::14
7100 Chassis(rw-config-dhcp-v6-pool) ->nis-domain-name myNisEnterprise.com
7100 Chassis(rw-config-dhcp-v6-pool) ->nis-dns-server 1111::12 1111::13 1111::14
7100 Chassis(rw-config-dhcp-v6-pool) ->nisp-domain-name myNispEnterprise.com
7100 Chassis(rw-config-dhcp-v6-pool) ->nis-dns-server 1111::12 1111::13 1111::14
7100 Chassis(rw-config-dhcp-v6-pool) ->sip-domain-name mySipEnterprise.com
7100 Chassis(rw-config-dhcp-v6-pool) ->sip-dns-server 1111::12 1111::13 1111::14
7100 Chassis(rw-config-dhcp-v6-pool) ->sntp-server 1111::15
7100 Chassis(rw-config-dhcp-v6-pool) ->unicast-server 1111::15
7100 Chassis(rw-config-dhcp-v6-pool) ->information-refresh 0 12 0
7100 Chassis(rw-config-dhcp-v6-pool) ->exit
7100 Chassis(rw-config) ->interface vlan 44
7100 Chassis(rw-config-intf-vlan.0.44) ->ipv6 address 9999::1/64
7100 Chassis(rw-config-intf-vlan.0.44) ->ipv6 dhcp server dhcpPool1
7100 Chassis(rw-config-intf-vlan.0.44) ->ipv6 forwarding
7100 Chassis(rw-config-intf-vlan.0.44) ->no shutdown
7100 Chassis(rw-config-intf-vlan.0.44) ->exit
7100 Chassis(rw-config) ->
```

## IPv6 DHCP Relay Source and Destination Interfaces

By default the router interface that receives the DHCP request is the DHCP request source interface for the router. You can configure an always-up global source interface for the device such as a loop-back interface. The configured global source-interface can be overridden at the interface level.

Use the **ipv6 dhcp relay source-interface** command in global configuration mode to configure a global source interface for the router.

Use the **ipv6 dhcp relay source-interface** command in interface configuration mode to have the specified interface override the globally configured source interface for this interface.

The destination server interface must be specified when the DHCPv6 destination server address is either link-local or multicast IPv6. Specifying a destination interface is not required if the DHCPv6 destination server address is a global address.

The DHCP Solicit message is a multicast message to the all DHCP server address (ff02::1:2). The all DHCP server address only crosses network segments when explicitly routed. If your network has multiple segments, you must configure a DHCP relay agent on the router interface for each segment, so that all DHCP solicit messages can be forwarded to your DHCP server.

If a destination interface is not specified, because the DHCPv6 server address is a global address, the interface is determined by a standard routing table lookup.

Use the **ipv6 dhcp relay destination** command in interface configuration mode to configure the IPv6 DHCP relay agent to forward an IPv6 DHCP request from a client or other relay agent to the destination server or next relay agent address.

## Configuring DHCP

This section provides details for the configuration of DHCP on the 7100-Series products.

[Table 9-11](#) lists DHCP parameters and their default values.

**Table 9-11 Default DHCP Parameters**

Parameter	Description	Default Value
DHCP interface state	Specifies whether DHCP is enabled or disabled on a routing interface.	disabled
number of ping packets	Specifies the number of packets a DHCP server sends to an IP address before assigning the address to a requesting client.	2
ping timeout	Specifies the amount of time the DHCP server will wait for a ping reply from an IP address before timing out.	500 milliseconds
information refresh	Specifies the amount of time an DHCPv6 client will wait before requesting an option information refresh from the DHCPv6 server.	1 day

[Procedure 9-10](#) describes enabling the DHCP feature and client configuration.

### Procedure 9-10 Enabling the DHCP Server and Configuring Automatic Address Assignment

Step	Task	Command(s)
1.	Enable DHCP on the routing interface in interface configuration command mode. DHCP is enabled by default.	<b>ip dhcp server</b>
2.	Configure the local address pool to be used as a DHCP subnet for automatic IP address assignment.	<b>ip local pool name subnet mask</b>
3.	Optionally, in local address pool configuration mode, exclude a range of IP addresses from the configured local pool subnet, specifying the beginning IP address and the number of additional addresses to exclude.	<b>exclude ip-address number</b>
4.	Enter DHCP address pool configuration command mode for the specified pool.	<b>ip dhcp pool name</b>
5.	Specify, in DHCP pool or client-class mode, the lease duration for an IP address dynamically assigned by a DHCP server to a client.	<b>lease {days [hours] [minutes]}</b>
6.	In DHCP pool configuration mode, enable DHCP host configuration mode and optionally associate a client class with a DHCP client.	<b>client-identifier unique-identifier [client-class name]</b>
7.	In DHCP pool configuration mode, specify parameters for a new DHCP client address.	<b>hardware-address hardware-address [type]</b>
8.	Specify, in configuration command mode, the number of packets a DHCP Server sends to a pool address as part of a ping operation.	<b>ip dhcp ping packets number</b>



### Procedure 9-10 Enabling the DHCP Server and Configuring Automatic Address Assignment (continued)

Step	Task	Command(s)
9.	Specify, in configuration command mode, the number of milliseconds the DHCP server will wait for a ping reply from an IP address before timing out.	<b>ip dhcp ping timeout</b> <i>milliseconds</i>
10.	In either configuration command mode or interface configuration mode, specify an always up source interface of IPv6 DHCP relay forwarded messages.	<b>ipv6 dhcp relay source-interface</b> <i>interface</i>
11.	In interface command mode, configure the IPv6 DHCP relay agent to forward an IPv6 DHCP request from a client or other relay agent to the destination server or relay agent address.	<b>ipv6 dhcp relay destination</b> <i>ipv6-address [destination-interface] [global] [vrf vrf]</i>

Table 9-12 describes how to configure the router.

### Table 9-12 Configuring Static IP Address Assignment

Task	Command(s)
Optionally, configure static IP address assignment in DHCP host configuration command mode by specifying an host IP address and network mask for a static DHCP binding.  Use either the <b>hardware-address</b> or <b>client-identifier</b> command in DHCP pool configuration command mode to enter host configuration command mode.	<b>host address</b> [ <i>mask   prefix-length</i> ]

Procedure 9-11 describes DHCP client configuration.

### Procedure 9-11 DHCP Client Configuration

Step	Task	Command(s)
1.	Optionally, in DHCP host or pool configuration command mode, specify a domain name for the DHCP client.	<b>domain-name</b> <i>name</i>
2.	Specify, in DHCP host or pool configuration command mode, one or more DNS server IP addresses to the DHCP clients.	<b>dns-server</b> <i>address [address2...address8]</i>
3.	Specify, in DHCP host or pool configuration command mode, one or more NetBIOS WINS servers to the DHCP clients.	<b>netbios-name-server</b> <i>address [address2...address8]</i>
4.	Specify, in DHCP host or pool configuration command mode, one or more node types to the DHCP clients. <ul style="list-style-type: none"> <li>• <b>h-node</b> — hybrid (recommended)</li> <li>• <b>b-node</b> — broadcast</li> <li>• <b>p-node</b> — peer-to-peer</li> <li>• <b>m-mode</b> — mixed</li> </ul>	<b>netbios-node-type</b> <i>type</i>

**Procedure 9-11 DHCP Client Configuration (continued)**

Step	Task	Command(s)
5.	Optionally, in DHCP host or pool configuration command mode, assign routers to a DHCP client's default router list.	<b>default-router</b> <i>address</i> [ <i>address2...address8</i> ]
6.	Specify, in DHCP host or pool configuration command mode, the default boot image for the DHCP client.	<b>bootfile</b> <i>filename</i>
7.	Optionally, in DHCP host or pool configuration command mode, specify the next server in the DHCP server boot process.	<b>next-server</b> <i>ip-address</i>
8.	Optionally, in DHCP host or pool configuration command mode, configure DHCP options.	<b>option code</b> [ <i>instance number</i> ] { <i>ascii string</i>   <i>hex string</i>   <i>ip address</i> }
9.	Optionally, in client configuration command mode, assign a name to a DHCP client. Optionally, assign the named client to a client class.	<b>client-name</b> <i>name</i> [ <b>client-class</b> <i>name</i> ]
10.	Optionally, in DHCP host or pool configuration command mode, configure a client class.	<b>client-class</b> <i>name</i>

[Procedure 9-12](#) describes enabling the DHCP feature and client configuration.

**Procedure 9-12 Configuring DHCPv6 Information Option Pools**

Step	Task	Command(s)
1.	In global configuration mode, create a DHCPv6 option pool and enter DHCPv6 pool configuration mode.	<b>ipv6 dhcp pool</b> <i>poolname</i>
2.	In interface configuration mode, assign a DHCPv6 information option pool to a DHCPv6 server.	<b>ipv6 dhcp server</b> <i>poolname</i>

[Procedure 9-13](#) describes DHCPv6 client configuration. All commands are configured in DHCPv6 pool configuration command mode.

**Procedure 9-13 DHCPv6 Client Configuration**

Step	Task	Command(s)
1.	Specify a domain name for the DHCPv6 client.	<b>domain-name</b> <i>name</i>
2.	Specify one or more DNS server IP addresses to the DHCP clients.	<b>dns-server</b> <i>address</i> [ <i>address2...address8</i> ]
3.	Specify one or more Network Information Services (NIS) domain names to return when responding to a DHCPv6 client request.	<b>nis-domain-name</b> <i>domain</i> [ <i>domain2 ... domain8</i> ]
4.	Specify one or more Network Information Services (NIS) servers to assign to DHCPv6 clients.	<b>nis-server</b> <i>address</i> [ <i>address2...address8</i> ]
5.	Specify one or more Network Information Services (NIS) version 2 domain names to return when responding to a DHCPv6 client request.	<b>nisv2-domain-name</b> <i>domain</i> [ <i>domain2 ... domain8</i> ]

**Procedure 9-13 DHCPv6 Client Configuration (continued)**

Step	Task	Command(s)
6.	Specify one or more Network Information Services (NIS) version 2 servers to assign to DHCPv6 clients.	<b>nis-server</b> <i>address</i> [ <i>address2...address8</i> ]
7.	Specify one or more Session Initiation Protocol (SIP) domain names to return when responding to a DHCPv6 client request.	<b>sip-domain-name</b> <i>domain</i> [ <i>domain2 ... domain8</i> ]
8.	Specify one or more Session Initiation Protocol (SIP) servers to assign to DHCPv6 clients.	<b>sip-server</b> <i>address</i> [ <i>address2...address8</i> ]
9.	Specify a Simple Network Time Protocol (SNTP) server to assign to DHCPv6 clients.	<b>sntp-server</b> <i>address</i>
10.	Specify a unicast server to assign to DHCPv6 clients.	<b>unicast-server</b> <i>address</i>
11.	Specify the amount of time a client should wait before refreshing information from the DHCPv6 server.	<b>information-refresh</b> { <i>infinite</i>   <i>days</i> [[ <i>hours</i> ] [ <i>minutes</i> ]]}

[Table 9-13](#) describes how to manage and display DHCP.

**Table 9-13 Managing and Displaying DHCP**

Task	Command(s)
To display IP DHCP bindings, in any command mode enter:	<b>show ip dhcp binding</b> [ <i>ip-address</i> ]
To display DHCP server statistics, in any command mode enter:	<b>show ip dhcp server statistics</b>
To delete one or all automatic DHCP address bindings, in configuration command mode enter:	<b>clear ip dhcp binding</b> { <i>address</i>   *}
To clear ip dhcp server statistics, in configuration command mode enter:	<b>clear ip dhcp server statistics</b>

## Node Alias Overview

Node alias provides for the defining of objects which can be used for the discovery of end systems on a per port basis. Because the 7100-Series firmware sees all packets that transit a port as members of a flow, node alias uses that flow defining capability to map key system objects such as VLAN ID, Source IP address, MAC address, host name, and protocol that define the end-users transiting the node alias enabled port. Enabling all ports for node alias allows for the building of a network wide cross-reference of key user elements providing the network administrator with a powerful troubleshooting tool.

Node alias creates an entry for each unique set of elements discovered when investigating the packets that transit the node alias enabled port. Node alias entries can be configured for all protocols or per protocol.

## Configuring Node Alias

This section describes how to configure Node Alias on the 7100-Series products.

[Procedure 9-14](#) describes how to configure node alias on switch ports.

**Procedure 9-14 Configuring Node Alias**

Step	Task	Command(s)
1.	Optionally disable node alias on switch ports. All ports and LAGs are enabled by default.	<b>set nodealias disable</b> [ <b>protocols protocols</b> ] port-string
2.	Optionally change the maximum number of entries allowed for the specified switch port.	<b>set nodealias maxentries</b> port-string

Table 9-14 describes how to display and manage the node alias on the 7100-Series device.

**Table 9-14 Managing Node Alias**

Task	Command(s)
To display the current port node alias state and maximum entries settings.	<b>show nodealias config</b> [port-string]
To display node alias entries for all or the specified port(s).	<b>show nodealias</b> [port-string]
To display node alias entries for the specified MAC address, optionally narrowing the search by protocol and port. The MAC address can be specified as a partial MAC address.	<b>show nodealias mac</b> mac_address [protocol] [port-string]
To display node alias entries for the specified protocol, optionally narrowing the search by port. In the case of the IP protocol, an IP address in full or partial form can be specified.	<b>show nodealias protocol</b> {protocol} [ip_address ip-address] [port-string]
To clear a specified node alias entry or all entries for the specified port(s).	<b>clear nodealias</b> {port port-string   alias-id alias-id   [protocols protocols]}
To reset node alias state to enabled and clear the maximum entries value for the specified port(s).	<b>clear nodealias config</b> port-string

## Setting Node Alias State and Max Entries

Node alias state and maximum entries settings are set using the **set nodealias** command in any command mode. Use the **show nodealias config** command to display the current nodealias state and maximum entries setting for this device.

The following example enables node alias on port ge.1.1, sets the maximum entries for ge.1.1 to 100, and displays all entries using the VRRP protocol:

```
7100-Series(rw)->set nodealias enable ge.1.1
7100-Series(rw)->set nodealias maxentries 100
7100-Series(rw)->show nodealias protocol vrrp ge.1.1
```

```
Port: ge.1.1    Time:    2009-07-24 16:20:37
```

```
-----
Alias ID       = 194020           Active           = true
Vlan ID        = 1                MAC Address      = 00-00-5e-00-01-01
Protocol       = vrrp           Rtr ID          = 0x01
Rtr priority   = 0xff
```

The following example displays all entries on port ge.1.1 with a MAC address beginning with 00-90:

```
7100-Series(rw)->show nodealias mac 00-90 ge.1.1
```

```
Port: ge.1.1 Time: 2009-07-24 16:28:47
```

```
-----
Alias ID      = 194067          Active          = true
Vlan ID      = 1           MAC Address     = 00-90-27-17-13-e7
Protocol     = ip          Source IP       = 10.21.2.95
```

The following example displays all entries on port ge.1.1 with an IP subnet of 10.21.\*.\*

```
7100-Series(rw)->show nodealias protocol ip ip_address 10.21 ge.1.1
```

```
Port: ge.1.1 Time: 2009-07-25 08:12:33
```

```
-----
Alias ID      = 194426          Active          = true
Vlan ID      = 1           MAC Address     = 00-00-5e-00-01-01
Protocol     = ip          Source IP       = 10.21.64.1
.
.
.
```

```
Port: ge.1.1 Time: 2009-07-25 08:25:15
```

```
-----
Alias ID      = 194460          Active          = true
Vlan ID      = 1           MAC Address     = 00-01-f4-5b-5f-a7
Protocol     = ip          Source IP       = 10.21.64.1
```

```
Port: ge.1.1 Time: 2009-07-25 08:14:45
```

```
-----
Alias ID      = 194435          Active          = true
Vlan ID      = 1           MAC Address     = 00-e0-63-86-2b-bf
Protocol     = ip          Source IP       = 10.21.64.2
```

## MAC Address Settings Overview

MAC address settings configuration provides for the ability to:

- Configure a timeout period for aging learned MAC addresses
- Limit specified layer two multicast addresses to specific ports within a VLAN
- Statically enter unicast MAC addresses into the filtering database (FID). Static MAC addresses can be permanent or ageable
- Enable the ability to treat static unicast MAC addresses as a multicast address

## Age Time

Both learned and statically configured MAC addresses can be assigned an age in seconds after which they will be flushed from the FID. The default value is 300 seconds.

Use the **set mac agetime** command in any command mode to configure the MAC age-time for MAC addresses on this device.

The following example sets the age-time for MAC addresses on this device to 600 seconds:

```
7100 Chassis(rw)->set mac agetime 600
7100 Chassis(rw)->show mac agetime
Aging time: 600 seconds
7100 Chassis(rw)->
```

## Multicast MAC Address VLAN Port Limit

Specified layer two multicast MAC addresses can be limited to specific ports within a VLAN. You can append or clear ports from the list of ports the multicast MAC address is dynamically learned on or flooded to.

Use the **set mac multicast** command in any command mode to limit the specified multicast MAC address to specific ports within a VLAN. This command creates a static MAC address that forces frames with a specific multicast destination address to be hardware switched on the specified VLAN. This command can also be used to flood a unicast MAC address. The command takes a list of ports. The specified port list scopes the flooding to a port set smaller than the egress port set for the VLAN. Multicast frames matching a static MAC address entry are transmitted on each of the specified ports (or all ports if no port is specified) that egresses the VLAN.



**Note:** A key purpose of the static multicast MAC address feature is to support the multicast version of various NIC-based proprietary load balancing technologies, including Network Load Balancing (NLB). See “[Network Load Balanced \(NLB\) Servers Configured for Multicast](#)” on page 9-37.

The following example specifies that multicast MAC address 00:a4:01:ff:0e:00:01 be limited to port ge.1.1 on VLAN 100:

```
7100 Chassis(rw)->set mac multicast 00:a4:01:ff:0e:01 100 ge.1.1
Warning: Unicast address converted to multicast 01-A4-01-FF-0E-01
```

Unicast MAC addresses can be statically entered into a FID for a single port. This entry can be configured as either permanent or ageable. If ageable, it will age out the same as a dynamically learned MAC address.

## Network Load Balanced (NLB) Servers Configured for Multicast

Network load balancer or similar proprietary load balancing technologies, comprised of multiple physical machines responding to a single “virtual” IP address, expect the switch to flood its traffic to all ports on the destination VLAN. The flooded traffic uses the Extreme Networks device soft forwarding path, subject to its rate limiters, instead of the device hardware forwarding path. This traffic will also compete for the slow path resources and the first packets from other new flows.

To force the virtual server packets to take a hardware switch path, configure a MAC address static entry in the Filter Database (FDB). To optimally support NLB, the VLAN with the NLB servers should consist of only ports that are attached to the NLB servers. non-NLB servers on this VLAN will receive the multicast frames that are typically only meant for the NLB servers.

Configure the device for enhanced routing functionality by entering the **set limit resource\_profile router1** command.

## Static MAC Address Entry

Use the **set mac unicast** command in any command mode to statically enter a unicast MAC address into a FID for a single port.

The following example statically enters unicast MAC address 00:a4:01:ff:0e:01 into FID 1 for port ge.1.1 and sets the MAC address to ageable:

```
7100 Chassis(rw)->set mac unicast 00:a4:01:ff:0e:01 1 ge.1.1 ageable
7100 Chassis(rw)->show mac fid 1
```

MAC Address	FID	Port	Type	Status
00-00-5E-00-01-01	1	ge.1.1	learned	
00-16-41-A8-8F-D8	1	ge.1.1	learned	
00-A0-C9-0A-8F-52	1	ge.1.1	learned	
00-A4-01-FF-0E-01	1	ge.1.1	mgmt	ageable
00-B0-D0-B7-D2-C5	1	ge.1.1	learned	

```
7100 Chassis(rw)->
```

## Unicast as Multicast

The unicast as multicast feature causes unicast searches in the filter data base to match on statically configured multicast entries using hardware forwarding. The unicast as multicast feature is used when a data stream originates from or is forwarded to a unicast address that then forwards it to multiple hosts, such as when using Network Load Balancing (NLB). When unicast as multicast is enabled on the device, a lookup is performed to determine if the unicast address has also been configured for multicast on the device. If a multicast address is found, packets are hardware forwarded out the configured VLAN and port(s) as defined in the static multicast configuration by extending the search phase of the Layer 2 lookup to match an unlearned destination MAC address against static multicast MAC entries. The unicast as multicast feature is configured by:

1. Using the **set mac multicast** command, in any command mode, to specify the MAC address to be treated as a multicast address, specifying the VLAN and egress port(s) to use
2. Using the **set mac unicast-as-multicast** command, in any command mode, to enable static unicast MAC addresses to be treated as multicast addresses on this device

The following command enables the unicast as multicast feature on this device:

```
7100 Chassis(rw)->set mac unicast-as-multicast enable
7100 Chassis(rw)->show mac unicast-as-multicast
Unicast as multicast: enabled
7100 Chassis(rw)->
```

## New and Moved MAC Address Detection

You can configure this device such that SNMP trap messaging is enabled globally or per port to send notifications, when a new MAC address is first detected, or a preexisting MAC address is moved.

Use the **set newaddrtrap** command in any command mode to enable SNMP trap messaging to report the detection of a new MAC address for the affected ports. Enabling SNMP trap messaging to report the detection of a new MAC address must be enabled globally and enabled on the

affected ports as two separate CLI entries. The new MAC address trap feature is disabled by default.

The following example configures SNMP trap messaging to send a notification when a new MAC address is detected on port ge.1.1:

```
7100 Chassis(rw)->set newaddrtrap enable
7100 Chassis(rw)->set newaddrtrap ge.1.1 enable
```

Use the **set movedaddrtrap** command in any command mode to enable SNMP trap messaging to report detection of a moved MAC address for the affected ports. Enabling SNMP trap messaging to report detection of a moved MAC address must be enabled globally and on the affected ports as two separate CLI entries. The moved MAC address trap feature is disabled by default.

The following example configures SNMP trap messaging to send a notification when a moved MAC address is detected on port ge.1.1:

```
7100 Chassis(rw)->set movedaddrtrap enable
7100 Chassis(rw)->set movedaddrtrap ge.1.1 enable
```

[Procedure 9-15](#) describes how to configure MAC address settings. All commands for this feature can be set in any command mode.

### Procedure 9-15 Configuring MAC Address Settings

Step	Task	Command(s)
1.	Optionally, change the age time for MAC addresses FID entries for this device.	<b>set mac agetime</b> <i>time</i>
2.	Optionally, limit a multicast MAC address to a specific port within a VLAN.	<b>set mac multicast</b> <i>mac-address vlan-id [port-string]</i> { <b>append</b>   <b>clear</b> }
3.	Optionally, enter a static unicast MAC address into the FID.	<b>set mac unicast</b> <i>mac-address fid receive-port [ageable]</i>
4.	Optionally, enable unicast MAC addresses to be treated as multicast MAC addresses on this device.	<b>set mac unicast-as-multicast</b> { <b>enable</b>   <b>disable</b> }
5.	Optionally, enable SNMP trap messaging to report the detection of new MAC addresses for the specified port by first globally enabling the feature, followed by enabling the affected ports.	<b>set newaddrtrap</b> { <b>enable</b>   <b>disable</b> } (Global) <b>set newaddrtrap</b> <i>port-string</i> { <b>enable</b>   <b>disable</b> } (Port)
6.	Optionally, enable SNMP trap messaging to report the detection of a moved MAC address for the specified port by first globally enabling the feature, followed by enabling the affected ports.	<b>set movedaddrtrap</b> { <b>enable</b>   <b>disable</b> } (Global) <b>set movedaddrtrap</b> <i>port-string</i> { <b>enable</b>   <b>disable</b> } (Port)

## Terms and Definitions

[Table 9-15](#) lists terms and definitions used in this system configuration discussion.

**Table 9-15 System Configuration Terms and Definitions**

Term	Definition
age time	The amount of time a non-permanent MAC address will stay in the FIB before becoming marked as invalid.
automatic address assignment	DHCP automatically assigns an IP address from a range of configured addresses to a client for a limited period of time



**Table 9-15 System Configuration Terms and Definitions (continued)**

<b>Term</b>	<b>Definition</b>
broadcast listening	An SNTP operational mode for which the SNTP server broadcasts the time adding a configured propagation delay value to compensate for the travel time of the packet from the SNTP server to the SNTP client.
Domain Name Server (DNS) resolver	A session layer protocol that maps network host names to IP addresses and vice versa.
Dynamic Host Configuration Protocol (DHCP)	A network layer protocol that implements automatic or manual assignment of IP addresses and other configuration information to client devices by servers.
entry	A grouping of key packet objects reported by node alias that define a single flow for this port.
FID	The filtering database that contains the MAC addresses for this device.
manual address assignment	The client's IP address is assigned by the network administrator, DHCP is used only to convey the assigned address to the client.
node alias	An 7100-Series feature that analyzes flows transiting a port for key packet objects that can be used as a cross-reference that port's end users.
poll-interval	The time between SNTP update requests by the client to the server in unicast operations mode.
poll-timeout	The time a unicast SNTP client waits before sending another update request to the SNTP server.
precedence	A value used to determine the order in which SNTP servers will be polled in unicast operational mode.
Secure Shell (SSH)	security feature provides a secure encrypted communications method between a client and the switch to the entire session, providing data privacy and integrity that is an alternative to the unsecure Telnet protocol.
Simple Network Time Protocol (SNTP)	A protocol that provides for the synchronizing of system time for managed devices across a network.
unicast as multicast	A feature that treats a unicast MAC address as if it were a multicast MAC address by extending the search phase of layer 2 lookup to match the unlearned destination MAC address against the static Multicast MAC entries on this device.
unicast polling	An SNTP operational mode for which the client directly requests updates from the SNTP server.

## Security Mode Configuration

This chapter provides information about configuring and monitoring security modes on 7100-Series devices.

For information about...	Refer to page...
<a href="#">How to Use Security Mode in Your Network</a>	10-1
<a href="#">Implementing Security Mode</a>	10-6
<a href="#">Configuring Security Mode</a>	10-6
<a href="#">Security Mode Configuration Example</a>	10-7
<a href="#">Terms and Definitions</a>	10-7

### How to Use Security Mode in Your Network

There are three aspects to setting the security mode on your device:

- FIPS security which determines the authentication and encryption algorithms supported on your system
- Security profile which provides for either normal (standard) operation or the setting of the government C2 security rating which, for a subset of security related commands, sets:
  - Non-standard command parameter default values
  - Non-standard command parameter range values
  - Access to the command
- Boot menu access which enables or disables access to the boot menu



**Note:** Super-user administrative privilege is required to access security mode configuration for FIPS, security profile, and boot menu access.

### FIPS Security Mode

When enabled, FIPS security mode puts the switch into Federal Information Processing Standards (FIPS) mode. FIPS security mode is a mode where only FIPS approved authentication and encryption algorithms and methods are used. FIPS security mode defaults to disabled. FIPS security mode must be manually enabled using the **set security fips mode** command. The **show security fips mode** command displays the current FIPS security mode state for the device.

If FIPS security mode is enabled, only the SHA1 authentication algorithm is supported.



**Note:** Changing the FIPS security mode of the switch requires a system reset.

## Security Profile Mode

C2 security mode is a security rating established by the U.S. National Computer Security Center (NCSC) and specifies that a product passes the Department of Defense (DoD) Trusted Computer System Evaluation Criteria (TCSEC) tests. A product that meets the C2 security mode rating provides at least the minimum allowable levels of confidence demanded for government agencies and offices and other organizations that process classified or secure information.

Security profile mode determines access to some commands based upon user mode and the sensitivity of the command to security considerations. Security profile can also affect parameter range and default values for some commands.

Security profile mode is set using the **set security profile** command and can be set to **C2** or **normal**. The **show security profile** command displays the current security profile setting for the device.



**Note:** Changing the security profile mode of the switch requires a system reset.

## Boot Access Security Mode

On the 7100-Series device, the boot menu is called the System Image Loader menu. This menu provides for such boot related functionality as the ability to:

- Delete an image file
- Set the boot file image
- Start a ZMODEM download
- Display available image files or current boot image file
- Clear the persistent storage

By default you can gain access to this menu as the device is booting by pressing any key once you see the line:

```
Press any key to enter System Image Loader menu
```

Pressing any key places you at the System Image Loader prompt:

```
[System Image Loader]:
```

See [“Setting the Boot Firmware Image”](#) on page 3-8 for additional boot menu information.

Access to the boot menu can be enabled or disabled using the **set security boot-access** command. The **show security boot-access** command displays the current boot menu access setting for the device.

Disabling access to the boot menu affects all user privilege modes, including super-user.

## Security Profile Mode Default Parameter Setting Changes

Some command parameter default settings change when changing the security profile mode. [Table 10-1](#) details command parameter default setting changes per security profile mode.

**Table 10-1 Security Profile Mode Command Parameter Default Setting Changes**

Description	Command	Normal Default	C2 Default
Sets the time an idle console, SSH or Telnet CLI session will remain connected before being logged out. See <a href="#">“Using the CLI”</a> on page 2-1 for additional configuration information.	<b>set logout</b>	10 minutes	15 minutes
Sets the number of minutes to lockout the default admin super-user account after maximum login attempts. See <a href="#">“User Management Overview”</a> on page 9-5 for additional configuration information.	<b>set system lockout time</b>	0 minutes	8 minutes
Sets the minimum interval in minutes between password changes allowed for non-super-users. See <a href="#">“User Management Overview”</a> on page 9-5 for additional configuration information.	<b>set system password change-frequency</b>	0 never	1 day
Sets the number of inactive days before a non-super-user account is locked out. See <a href="#">“User Management Overview”</a> on page 9-5 for additional configuration information.	<b>set system lockout inactive</b>	0 never	90 days
Sets a grace period in either the number of logins or days before the password is locked out. See <a href="#">“User Management Overview”</a> on page 9-5 for additional configuration information.	<b>set system password grace-period logins</b>	0 off	3 logins
Sets the number of days after a password expires before the password is locked out. See <a href="#">“User Management Overview”</a> on page 9-5 for additional configuration information.	<b>set system password grace-period time</b>	0 off	30 days
Sets the SNMP user configuration privacy.	<b>set snmp user encryption</b>	usmNoPriv Protocol	usmAesCfg 128Protocol
Sets the SNMP user configuration authentication.	<b>set snmp user authentication</b>	usmNoAuth Protocol	usmHMACSHA AuthProtocol

## Security Profile Mode Parameter Range Changes

The `set system lockout attempts` command parameter range can change when changing the security profile mode. [Table 10-2](#) details the command parameter range change per security profile mode.

**Table 10-2 Security Profile Mode Command Parameter Range Changes**

Description	Command	Normal Default	C2 Default
Sets the number of failed login attempts before locking out (disabling) a read-write or read-only user account. See <a href="#">“User Management Overview”</a> on page 9-5 for additional configuration information.	<code>set system lockout attempts</code>	1 – 15	2 - 5

## C2 Security Profile Mode Command Access Changes

Some commands that are accessible in normal security profile mode are not accessible in C2 security profile mode. For some commands this change in access depends upon the user mode. [Table 10-3](#) details security profile mode command access changes.

**Table 10-3 Security Profile mode Command Access Changes**

Description	Command	Normal Default	C2 Default
Sets the authentication type required for this user as MD5 or SHA. Only MD5 is affected by C2 security profile mode.	<code>set snmp user authentication md5</code>	Allowed	MD5 Cloaked
Sets the privacy protocol to Advanced Encryption Standard (AES) or Data Encryption Standard (DES). Only DES is affected by C2 security profile mode.	<code>set snmp user encryption des</code>	Allowed	DES Cloaked
Creates a new SNMPv3 user.	<code>set snmp user</code>	RW and SU user modes	SU user mode only
Sets the properties for one or more console ports. Only VT100 is affected by C2 security profile. See <a href="#">“Console Port Parameters”</a> on page 6-3 for additional configuration information.	<code>set console vt100</code>	RW and SU user modes	SU user mode only

## C2 Security Profile Mode Read-Write User Mode Changes

Some Read-Write user mode functionality accessible in normal security profile mode is not accessible when in the C2 security profile mode. [Table 10-4](#) details Read-Write user mode functionality that is not accessible when in C2 security profile mode.

**Table 10-4 Read-Write Functionality Not Accessible in C2 Security Profile Mode**

Description	Command
Secure directory including secure logs. See <a href="#">“Configuration and Image File Display Commands”</a> on page 3-9 for additional configuration information.	<code>dir</code>

**Table 10-4 Read-Write Functionality Not Accessible in C2 Security Profile Mode**

Description	Command
The display of messages logged on all blades. See <a href="#">“Interpreting Messages”</a> on page 25-6 for additional configuration information.	<b>show logging buffer</b>
Script access to secure logs. See <a href="#">“Running a Configuration Script”</a> on page 3-8 for additional configuration information.	<b>script</b>
Display technical support-related information output. See <a href="#">Table 9-3</a> on page 9-3 for additional configuration information.	<b>show support</b>
Display of encrypted passwords is cloaked in the <b>show config</b> command output. See <a href="#">“Executing a Configuration”</a> on page 3-3 for additional configuration information.	<b>show config</b>
Display of encrypted passwords is cloaked in the <b>show file</b> command output. See <a href="#">Table 9-3</a> on page 9-3 for additional configuration information.	<b>show file</b>
The non-append version of the configure command is not available. See <a href="#">“Executing a Configuration”</a> on page 3-3 for additional configuration information.	<b>configure</b>
Ability to create, modify, or delete snmp users, access views, traps configuration and engine ID is not available. See <a href="#">“Configuring SNMP”</a> on page 17-7	<b>{set   clear} snmp {user   access   notify   engine-id}</b>
Ability to create, modify, or delete the authentication login method. See <a href="#">“Setting the Authentication Login Method”</a> on page 9-7 for additional configuration information.	<b>{set   clear} authentication login</b>
Ability to create, modify, or delete system login, lockout, or password. See <a href="#">“User Management Overview”</a> on page 9-5 for additional configuration information.	<b>{set   clear} system {login   lockout   password}</b>
Ability to create, modify, or delete console settings. See <a href="#">“Console Port Parameters”</a> on page 6-3 for additional configuration information.	<b>{set   clear} console</b>
Ability to create, modify, or delete logging local, application, default, server or here. See <a href="#">“Syslog Overview”</a> on page 25-2 for additional configuration information.	<b>{set   clear} logging {local   application   default   server   here}</b>
Ability to display or set the C2 security profile mode.	<b>{show   set} security profile</b>
Ability to display or set the FIPS security mode.	<b>{show   set} security fips mode</b>
Ability to display or set the security boot access mode.	<b>{show   set} security boot-access</b>
Ability to clear the configuration on all modules.	<b>clear config all</b>

## C2 Security Profile Mode Read-Only User Mode Changes

Some Read-Only user mode functionality accessible in normal security profile mode is not accessible when in the C2 security profile mode. The following table provides a list of Read-Only user mode functions that are not accessible when in C2 security profile mode.

**Table 10-5 Read-Only Functionality Not Accessible in C2 Security Profile Mode**

Description	Command
Secure directory including secure logs.	<b>dir</b>
The display of messages logged on all blades.	<b>show logging buffer</b>
Script access to secure logs.	<b>script</b>
Display technical support-related information output.	<b>show support</b>
Display of encrypted passwords is cloaked in the <b>show config</b> and <b>show file</b> command outputs.	<b>show config</b> <b>show file</b>
The non-append version of the configure command is not available.	<b>configure</b>
Ability to display or set the C2 security profile.	<b>show security profile</b>
Ability to display the FIPS security mode.	<b>show security fips mode</b>
Ability to display the security boot access mode.	<b>show security boot-access</b>

## Implementing Security Mode

To implement security mode on your network:

- Optionally restrict authentication and encryption algorithm support on the device to FIPS approved algorithms by enabling FIPS mode.
- Optionally set the security profile for the device to the C2 security level by setting the security profile to C2 mode.
- Optionally disable access to the system boot (System Image Loader) menu.

## Configuring Security Mode

Table 10-6 describes security mode configuration on the Extreme Networks 7100-Series devices.

**Table 10-6 Configuring Security Mode on the Device**

Task	Command(s)
Optionally, enable FIPS security mode restricting authentication and encryption algorithms to FIPS approved algorithms.	<b>set security fips mode {enable   disable}</b>
Optionally, set the security profile to C2 mode, changing a subset of security sensitive command default and range values, as well as command access, to meet the C2 security rating specification.	<b>set security profile {c2   normal}</b>
Optionally, disable access to the boot menu during bootup.	<b>set security boot-access {enable   disable}</b>

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

## Security Mode Display Commands

Table 10-7 lists security mode show commands.

**Table 10-7 Security Mode Show Commands**

Task	Command
To display the current boot access state for this device.	<b>show security boot-access</b>
To display the current security FIPS mode state for the device.	<b>show security fips mode</b>
To display the current security profile for the device.	<b>show security profile</b>

Refer to the *Extreme Networks 7100-Series CLI Reference* for a description of the output of each command.

## Security Mode Configuration Example

This security mode configuration example:

- Enables FIPS mode on the device, restricting the authentication and encryption algorithms to the FIPS approved SHA1 algorithm
- Sets the device security profile to C2:
  - Changing command parameter defaults as specified in Table 10-1 on page 10-3
  - Changing command parameter ranges as specified in Table 10-2 on page 10-4
  - Changing command access as specified in Table 10-3 on page 10-4
  - Denying read-write access to commands specified in Table 10-4 on page 10-4
  - Denying read-only access to commands specified in Table 10-5 on page 10-6
- Disables access to the boot menu during bootup on the device

```
7100 Chassis(su)->set security fips mode enable
This command will reset the system. Are you sure you want to continue? (y/n) [n]y
Resetting system ...
7100 Chassis(su)->set security profile c2
This command will reset the system. Are you sure you want to continue? (y/n) [n]y
Resetting system ...
7100 Chassis(su)->set security boot-access disable
```

## Terms and Definitions

Table 10-8 lists terms and definitions used in this security mode configuration discussion.

**Table 10-8 Security Mode Configuration Terms and Definitions**

Term	Definition
C2 security mode	C2 security mode is a security rating established by the U.S. National Computer Security Center (NCSC) and specifies that a product passes the Department of Defense (DoD) Trusted Computer System Evaluation Criteria (TCSEC) tests.



**Table 10-8 Security Mode Configuration Terms and Definitions (continued)**

<b>Term</b>	<b>Definition</b>
FIPS security	FIPS security mode puts the switch into Federal Information Processing Standards (FIPS) mode, where only FIPS approved authentication and encryption algorithms and methods are used.
security profile	Security profile sets the device security mode to either a normal (standard) level of security or to the C2 security mode.
system boot menu	The System Image Loader menu accessible during boot up of an 7100-Series device.

---

## IPsec Protocol Configuration

This chapter provides information about configuring and monitoring the IPsec protocol on 7100-Series devices.

For information about...	Refer to page...
<a href="#">How to Use IPsec in Your Network</a>	11-1
<a href="#">IPsec Implementation Requirements</a>	11-2
<a href="#">Understanding the IPsec Protocol</a>	11-3
<a href="#">Configuring IPsec</a>	11-7
<a href="#">Terms and Definitions</a>	11-12



**Note:** The IPsec protocol requires an advanced router license. See “[License Overview](#)” on page 9-10 for 7100-Series advanced router license details.

### How to Use IPsec in Your Network

The Internet Protocol Security Architecture (IPsec), defined in RFC 4301, describes how to provide a set of security services for traffic at the IP layer in both IPv4 and IPv6 environments. As described in the RFC, for this release, security services are provided through use of the Encapsulating Security Payload (ESP) traffic security protocol, and through the use of cryptographic key management procedures and protocols.

The IPsec implementation on the 7100-Series provides the following functionality:

- IPsec and IKE (Internet Key Exchange protocol) are defined for the RADIUS host application only. This implementation supports configuring the default Security Association (SA) with servers configured for RADIUS, and the RADIUS application helps define the IPsec flow.
- Only the Encapsulating Security Payload (ESP) mode of operation is supported. Authentication Header (AH) mode is not supported.
- IKEv1 is supported.



**Note:** Although the use of certificates will be supported for IPsec in future releases, in the current release, only use of a shared secret is supported.

- HMAC-SHA1 is the supported IKE integrity mechanism.
- 3DES and the Advanced Encryption Standard (AES) encryption algorithms are supported. AES supports key lengths of 128, 192, and 256 bits.

- IPsec does not prevent the independent simultaneous use of MSCHAP-V2 style encryption of user passwords between the switch and the RADIUS server.
- If FIPS security mode is enabled, using the **set security fips mode** command, only the SHA1 authentication algorithm is supported.

## IPsec Implementation Requirements

The following parameter configuration is required to implement IPsec on your network:

- IPsec is disabled by default. You must enable IPsec.
- IPsec is disabled by default for RADIUS transactions. You must configure the RADIUS management MS-CHAPv2 password attribute for IPsec to work. See the **set radius mgmt attribute** command details in the RADIUS commands chapter of the *Extreme Networks 7100-Series CLI Reference* for information on configuring the RADIUS management MS-CHAPv2 password attribute.

Configure the IPsec default instance by assigning an IKE map to it.

Optional IPsec and IKE configuration includes enabling IPsec traps.

## Required Manual Configuration

For this release, a number of IPsec parameters do not support default values and must be manually configured.

### IKE Proposal

For this release, the following IKE proposal parameters do not support default values and must be manually configured:

- The IKE Diffie-Hellman group.
- The IKE proposal encryption.
- The IKE proposal SHA1 hash and integrity (authentication).

### IKE Policy

The following IKE policy parameters must be manually configured:

- The Authentication Pre-Shared Key (PSK)
- The IKE policy lifetime. This release does not provide a default value.
- The policy SA peer (server).
- The IKE proposal assignment to the IKE policy.
- The IKE version. This release does not provide a default value.

### IKE Map

For this release, the following IKE map parameters must be manually configured:

- The UDP protocol. This release does not provide a default value.
- The destination IPv4 or IPv6 address (server).
- The encapsulation mode. This release does not provide a default value.
- The IKE map lifetime and bandwidth. This release does not provide a default value.
- The IKE policy assigned to the IKE map.

- The IKE proposal assigned to the IKE map.
- The source IPv4 or IPv6 address (local device).

## Understanding the IPsec Protocol

IPsec is an end-to-end security scheme protocol suite that secures IP communications using authentication and encryption of each communication session IP packet. IPsec can be used to protect data flows on a host-to-host and host-to-network basis. IPsec protects any application traffic across an IP network. Applications do not need to be specifically designed to use IPsec.

The 7100-Series IPsec implementation uses the ESP and SA protocols from the IPsec protocol suite. ESP provides for packet:

- Authenticity – ensures that the owner of the packet is who he claims to be
- Integrity – ensures that the contents of the packet have not been tampered with
- Confidentiality – ensures that information is accessible only to those authorized to have access

ESP operates directly on top of IP, using IP protocol number 50.

The Security Association (SA) protocol provides a bundle of algorithms and data required for ESP operations that are the basis for IPsec. The algorithms and data configured within an SA are used to encrypt and authenticate a particular flow in one direction. In a standard bi-directional communications session, two SAs are used, one for each direction. Security associations are established using the Internet Security Association and Key Management Protocol which provides for manual configuration of pre-shared secrets (keys) using the Internet Key Exchange (IKE).

IPsec identifies the SA that determines the protection to provide to an outgoing packet based upon a Security Parameter Index (SPI) and the packet header destination address. The SPI is an index to the security association database.

The 7100-Series IPsec implementation supports the configuration of a default SA. A default SA is configured by entering the IPsec default instance configuration mode and assigning an IKE map to the default SA.

### IKE Map

An IKE map groups together all algorithms and parameters that make up the SA. An IKE map contains the following parameters:

- The IKE proposal which groups the IKE map algorithms configured for the SA
- The IKE policy which groups policy related parameters configured for the SA
- The source and destination IP address and port for the SA
- The encapsulation type for the SA
- The map lifetime in time and bandwidth
- The transmission protocol (UDP) used by the SA
- Whether or not encryption is required

Use the **crypto ike-map** command in global VRF router configuration mode to create or modify an IKE map and enter IKE map configuration mode.

### IKE Proposal

The IKE proposal groups together the IKE map algorithms configured for the SA.

There are two IKE modes to which proposals are assigned: main mode and quick mode. The same IKE proposal can be assigned to both modes, or each mode can be assigned a unique IKE proposal depending upon your configuration needs.

The main mode or key exchange proposal is assigned to an IKE map in IKE map configuration mode. Main mode is the IKE negotiation that establishes a secure channel, known as the Internet Security Association and Key Management Protocol (ISAKMP) SA, between two devices.

Quick mode (also known as Phase 2) is the IKE negotiation that establishes a secure channel between two computers to protect data. Quick mode negotiates on behalf of the IPsec SAs. During quick mode, keying material is refreshed or, if necessary, new keys are generated. The quick mode proposal is assigned to an IKE policy using the **proposal** command in IKE policy configuration mode.

Use the **crypto ike-proposal** command in global VRF router configuration mode to create or modify an IKE proposal. Specify the name of the IKE proposal when entering the command. Upon entering the command, you are placed in IKE proposal configuration mode for the named proposal.

See [Table 11-1](#) for a description of IKE proposal parameters.

Use the **proposal** command in IKE map configuration mode to assign a main mode (key exchange) proposal to an IKE map.

**Table 11-1 IKE Proposal Parameters**

Parameter	Description
<b>IKE Diffie-Hellman (DH) group</b>	<p>IKE Diffie-Hellman (DH) group is a key derivation algorithm that generates the IPsec SA key. There are three algorithms supporting key sizes 768, 1024, and 2048 bits. The larger the generated key, the greater the security, but also the greater the system overhead.</p> <p>Use the <b>dh_group</b> command in IKE proposal configuration mode to set the IKE DH group algorithm for the proposal.</p>
<b>Encryption</b>	<p>Encryption is the process of transforming information, usually referred to as plaintext, using an algorithm, called a cipher, to make it unreadable to anyone except those possessing the associated key. The IKE proposal supports four encryption types:</p> <ul style="list-style-type: none"> <li>• <b>3des</b> – Triple Data Encryption Standard encryption algorithm</li> <li>• <b>aes128cbc</b> – The Advanced Encryption Standard (AES) 128 bit key size Cipher-Block Chaining (CBC) encryption algorithm.</li> <li>• <b>aes192cbc</b> – The Advanced Encryption Standard (AES) 192 bit key size Cipher-Block Chaining (CBC) encryption algorithm.</li> <li>• <b>aes256cbc</b> – The Advanced Encryption Standard (AES) 1256 bit key size Cipher-Block Chaining (CBC) encryption algorithm.</li> </ul> <p>This release does not support a default encryption algorithm. You must manually enter an encryption algorithm. Use the <b>encryption</b> command in IKE proposal configuration mode to set the encryption algorithm for the IKE proposal.</p>
<b>Hash</b>	<p>The hash algorithm is used during phase 1 negotiation between the SA authenticating devices. This release supports the Secure Hash Algorithm 1 (SHA1) hash. This release does not support a hash default value. You must manually enter the hash algorithm for one to be configured.</p> <p>Use the <b>hash</b> command in IKE proposal configuration mode to configure the hash algorithm for the IKE proposal.</p>

**Table 11-1 IKE Proposal Parameters (continued)**

Parameter	Description
Integrity	<p>Integrity, also referred to as data authentication, verifies that the data has not been altered as opposed to a user authentication which verifies the identity of the user. This release supports SHA1 integrity. SHA1 produces a 160-bit message digest for which no known attacks or partial attacks have yet been demonstrated. This release does not support a default integrity algorithm. You must manually enter the integrity algorithm for one to be configured.</p> <p>Use the <b>integrity</b> command in IKE proposal configuration mode to configure the integrity algorithm for the IKE proposal.</p>

## IKE Policy

The IKE policy groups together policy related parameters configured for the SA. Use the **crypto ike-policy** command in global VRF router configuration mode to create or modify an IKE policy. Specify the name of the IKE policy when entering the command. Upon entering the command you are placed in IKE policy configuration mode for the named policy. See [Table 11-2](#) for a description of IKE policy parameters.

**Table 11-2 IKE Policy Parameters**

Parameter	Description
Authentication Pre-shared Key	<p>The authentication PSK is a pre-shared authentication key that is used to initiate the connection and exchange encryption keys during the session.</p> <p>Use the <b>authentication psk</b> command in IKE policy configuration mode to configure the authentication pre-shared key for the SA.</p>
Initial Contact	<p>If the local host has rebooted, peers may have SAs that are no longer valid. If the initial contact feature is enabled, upon reboot an initial contact message is sent to a peer so that it will delete old SAs.</p> <p>Use the <b>initial-contact</b> command to enable the initial contact feature for the SA. The initial contact feature is disabled by default.</p>
Lifetime	<p>The IKE policy lifetime specifies the life cycle of an ISAKMP SA and is configured in minutes. The policy lifetime determines when a policy times out. A lifetime renegotiation automatically occurs before the lifetime is to expire. If the renegotiation is unsuccessful, the policy expires.</p> <p>Use the <b>lifetime time</b> command in IKE policy configuration mode to configure an IKE policy timeout period.</p>
Passive Mode	<p>Passive mode configures the IKE policy to wait for the peer to initiate the IKE session. By default a device is in active mode and constantly polls to see if the peer is up.</p> <p>Use the <b>passive</b> command in IKE policy configuration mode to configure the IKE policy for passive mode.</p>

**Table 11-2 IKE Policy Parameters (continued)**

Parameter	Description
Peer	<p>An IPv4 or IPv6 peer is specified for the SA using the <b>peer</b> command in IKE policy configuration mode.</p> <p>The IKE policy peer configuration determines whether the associated map is IPv4 or IPv6 during the phase 1 main mode negotiation. The CLI allows you to enter an inconsistent configuration between the local and peer address IP types. Inconsistent local and peer configurations will cause the IKE map to not be programmed.</p>
Proposal (Quick Mode)	<p>The quick mode proposal, used to establish and refresh user-level SAs, is assigned to an IKE policy using the <b>proposal</b> command in IKE policy configuration mode.</p>
Version	<p>The 7100-Series supports IKE version 1 for this release. Use the <b>version</b> command in IKE policy configuration mode to specify the IKE version used for the policy.</p> <p>This release does not support a default IKE version. You must manually enter an IKE version.</p>

## Source and Destination Address and Port

Source and destinations addresses need to be configured for the IKE map. The source address is the local address. The destination address is the remote address. Both IPv4 and IPv6 addresses are supported. Address ranges are supported using the slash (/) length notation.

Use the **dst** command in IKE map configuration mode to configure a destination address for the IKE map.

Use the **src** command in IKE map configuration mode to configure a source address for the IKE map.

The default SA port for the IKE map is any port. The source and destination ports can be specified for the SA when a specific protocol such as UDP is being authenticated. If you wish to limit the IKE map to a specific port, configure the source and destination port.

Use the **dst-port** command in IKE map configuration mode to configure the IKE map for a specific destination port.

Use the **src-port** command in IKE map configuration mode to configure the IKE map for a specific source port.

## Encapsulation

The SA encapsulation is determined by the type of communications required and determines whether the whole packet or only the data portion of the packet is encrypted and authenticated. There are two modes of encapsulation:

- Transport mode is used for host-to-host communications. In transport mode, only the transferred data of the IP packet is encrypted or authenticated. The routing is intact, since the IP header is neither modified nor encrypted; however, when the authentication header is used, the IP addresses cannot be translated, because to do so would invalidate the hash value.
- Tunnel mode is used to create virtual private networks. In tunnel mode, the entire IP packet is encrypted or authenticated. It is then encapsulated into a new IP packet with a new IP header.

This release does not support a default SA encapsulation. You must manually configure IKE map encapsulation.

Use the **encapsulation** command in IKE map configuration mode to specify the encapsulation mode to use for the SA.

## SA Lifetime

A lifetime can be set for the SA in both seconds and aggregate bandwidth. When a lifetime expires the SA is renegotiated as a security measure.

This release does not support a default SA lifetime. You must manually configure an SA lifetime for this IKE map.

Use the **lifetime** command in IKE map configuration mode to configure a lifetime value for the SA.

## Transmission Protocol

The UDP protocol is supported for SA packet transmission. This release does not support a default SA transmission protocol. You must manually configure an SA transmission protocol for this IKE map.

Use the **protocol udp** command to specify UDP as the SA transmission protocol.

## Encryption Request

By default, encryption is required to be used for the SA both locally and by the peer. If the peer does not support encryption, packets are not sent for the SA. If encryption request is enabled and the peer does not support encryption, packets are sent unencrypted. Use the **request** command to set the requirement for encryption to request for the SA. Request is disabled by default.

# Configuring IPsec

## IKE Proposal Configuration

[Procedure 11-1](#) describes IKE proposal configuration on the Extreme Networks 7100-Series devices.

### Procedure 11-1 Configuring an IKE Proposal

Step	Task	Command(s)
1.	Enter IKE proposal configuration mode, from the global VRF router configuration mode, to create a new or modify an existing IKE proposal.	<b>crypto ike-proposal</b> <i>proposal-identifier</i>
2.	In IKE proposal configuration mode, configure the IKE Diffie-Hellman (DH) key exchange group for the SA: <ul style="list-style-type: none"> <li>• DH group 1 (modp768)</li> <li>• DH group 2 (modp1024)</li> <li>• DH group 14 (modp2048)</li> </ul>	<b>dh_group</b> {1   2   14}
3.	In IKE proposal configuration mode, configure the encryption algorithm for the IKE proposal.	<b>encryption</b> {3des   aes128cbc   aes192cbc   aes256cbc}



**Procedure 11-1 Configuring an IKE Proposal (continued)**

Step	Task	Command(s)
4.	In IKE proposal configuration mode, configure the hash algorithm for the IKE proposal.	<b>hash sha1</b>
5.	In IKE proposal configuration mode, configure the integrity (data authentication) algorithm for the IKE proposal.	<b>integrity sha1</b>

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

## IKE Policy Configuration

[Procedure 11-2](#) describes IKE policy configuration on the Extreme Networks 7100-Series devices.

**Procedure 11-2 Configuring an IKE Policy**

Step	Task	Command(s)
1.	Enter IKE policy configuration mode, from the global VRF router configuration mode, to create a new or modify an existing IKE policy.	<b>crypto ike-policy</b> <i>policy-identifier</i>
2.	In IKE policy configuration mode, configure the authentication pre-shared key (PSK) for the IKE policy.	<b>authentication psk</b> <i>pre-shared-key</i>
3.	In IKE policy configuration mode, optionally enable initial contact feature for the IKE policy.	<b>initial-contact</b>
4.	In IKE policy configuration mode, configure the lifetime for the IKE policy.	<b>lifetime time</b> <i>minutes</i>
5.	In IKE policy configuration mode, optionally enable passive mode for the IKE policy	<b>passive</b>
6.	In IKE policy configuration mode, configure the SA peer for the IKE policy.	<b>peer</b> <i>address</i>
7.	In IKE policy configuration mode, assign an IKE proposal to the IKE policy.	<b>proposal</b> <i>proposal-identifier</i>
8.	In IKE policy configuration mode, configure the IKE version for the IKE policy.	<b>version</b> <i>version</i>

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

## IKE Map Configuration

[Procedure 11-3](#) describes IKE map configuration on the Extreme Networks 7100-Series devices.

**Procedure 11-3 Configuring an IKE Map**

Step	Task	Command(s)
1.	Enter IKE map configuration mode, from the global VRF router configuration mode, to create a new or modify an existing IKE map.	<b>crypto ike-map</b> <i>map-identifier</i>
2.	In IKE map configuration mode, configure a source address for the IKE map.	<b>src</b> <i>address</i>

**Procedure 11-3 Configuring an IKE Map (continued)**

Step	Task	Command(s)
3.	In IKE map configuration mode, optionally configure a source port for the IKE map. The source port defaults to any port.	<b>src-port</b> <i>port</i>
4.	In IKE map configuration mode, configure a destination (peer) address for the IKE map.	<b>dst</b> <i>address</i>
5.	In IKE map configuration mode, optionally configure a destination port for the IKE map. The destination port defaults to any port.	<b>dst-port</b> <i>port</i>
6.	In IKE map configuration mode, configure the encapsulation mode for the IKE map.	<b>encapsulation</b> { <i>tunnel</i>   <i>transport</i> }
7.	In IKE map configuration mode, configure the lifetime in time or bandwidth for the IKE map. Both values can be configured using separate command entries.	<b>lifetime</b> { <i>time minutes</i>   <i>bandwidth kilobytes</i> }
8.	In IKE map configuration mode, assign the specified IKE policy to the IKE map.	<b>policy</b> <i>policy-identifier</i>
9.	In IKE map configuration mode, assign the specified main mode key exchange IKE proposal to the IKE map.	<b>proposal</b> <i>proposal-identifier</i>
10.	In IKE map configuration mode, configure the IKE map with the UDP transmission protocol.	<b>protocol</b> <i>udp</i>
11.	In IKE map configuration mode, request that encryption be used by the SA and to not used it if encryption is not supported by the peer. Defaults to encryption required.	<b>request</b>

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

## IPsec Configuration

[Procedure 11-4](#) describes IPsec configuration on the Extreme Networks 7100-Series devices.

**Procedure 11-4 Configuring IPsec**

Step	Task	Command(s)
1.	Enter the IPsec default SA configuration mode, from the global VRF router configuration mode, to configure IPsec on the device.	<b>crypto ipsec default</b>
2.	In IPsec default SA configuration mode, assign an IKE map to the IPsec default SA.	<b>ike map</b> <i>ike-map</i>
3.	In global VRF router configuration mode, optionally enable IPsec traps.	<b>crypto ipsec trap-enable</b>
4.	In global VRF router configuration mode, enable IPsec on the router.	<b>crypto ipsec enable</b>

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

## IPsec Display Commands

Table 11-3 lists IPsec show commands.

**Table 11-3 IPsec Show Commands**

Task	Command
To display IKE statistics.	<b>show ike stats</b>
To display IKE proposal configuration.	<b>show ike proposal</b>
To display IKE policy configuration.	<b>show ike policy</b>
To display IKE map configuration.	<b>show ike map</b>
To display IKE SA information.	<b>show ike sa</b>
To display IPsec counters.	<b>show ipsec counters</b> [all   ipsec   global   memory   resources   task]
To display IPsec map information.	<b>show ipsec map</b>
To display IPsec SA information.	<b>show ipsec sa</b> [spi spi] [instance_id instance_id] [index index] [ipv4   ipv6] [brief]
To display IPsec flow information.	<b>show ipsec flow</b> [spi spi] [instance_id instance_id] [index index] [ipv4   ipv6] [brief]
To display IPsec instance information.	<b>show ipsec instance</b> [vlan vlan-id] [instance_id instance_id] [index index] [static   dynamic] [brief]
To display IPsec interface information.	<b>show ipsec instance</b> [vlan vlan-id] [instance_id instance_id] [static   dynamic] [brief]

Refer to the *Extreme Networks 7100-Series CLI Reference* for a description of the output of each command.

## IPsec Configuration Example

No specific order is required for configuring the IKE map that is applied to the default SA. Any required parameter or algorithm not properly configured will display as an incomplete configuration, but will not prevent you from configuring other IKE map parameters. A suggested order of IPsec configuration is:

1. Configure the main mode key exchange IKE proposal for the SA. Optionally configure a quick mode data protection IKE proposal if the quick mode proposal is different from the main mode proposal.
2. Configure the IKE policy for the SA.
3. Configure the IKE map for the SA.
4. Enter IPsec default instance configuration mode.
5. Apply the IKE map to the default SA
6. Enable IPsec

The following configuration example will follow the IPsec configuration order suggested above.

### IKE Proposal

As indicated in “[IKE Proposal](#)” on page 11-3, there are two IKE modes to which an IKE proposal is assigned:

- **Main**, which is assigned to an IKE map
- **Quick**, which is assigned to an IKE policy

Each IKE mode can be assigned a unique IKE proposal or the same proposal may be assigned to both modes, depending upon your configuration requirements. For this example we will configure a single IKE proposal named **winRadiusPro**, to be used in both IKE modes, with the following values:

- IKE Diffie-Hellman key exchange group – **14**
- Encryption – **aes128cbc**
- Hash – **SHA1**
- Integrity – **SHA1**

```
7100 Chassis(su)->configure
7100 Chassis(su-config)->crypto ike-proposal winRadiusPro
7100 Chassis(su-crypto-proposal)->dh_group 14
7100 Chassis(su-crypto-proposal)->encryption aes128cbc
7100 Chassis(su-crypto-proposal)->hash sha1
7100 Chassis(su-crypto-proposal)->integrity sha1
7100 Chassis(su-crypto-proposal)->exit
7100 Chassis(su-config)->
```

### IKE Policy

The IKE policy for this example is named **winRadiusPol**. The initial contact and passive mode features will not be enabled for this configuration. The **winRadiusPol** IKE policy is configured with the following values:

- Authentication pre-shared key – **testkey**
- Lifetime – **360** minutes
- Peer address – **1.1.191.22**
- IKE quick proposal – **winRadiusPro**
- IKE version – **1**

```
7100 Chassis(su-config)->crypto ike-policy winRadiusPol
7100 Chassis(su-crypto-policy)->authentication psk testkey
7100 Chassis(su-crypto-policy)->lifetime time 360
7100 Chassis(su-crypto-policy)->peer 1.1.191.22
7100 Chassis(su-crypto-policy)->proposal winRadiusPro
7100 Chassis(su-crypto-policy)->version 1
7100 Chassis(su-crypto-policy)->exit
7100 Chassis(su-config)->
```

### IKE Map

The IKE map for this example is named **winRadius**. IKE map parameters are configured with the following values:

- IKE main proposal – **winRadiusPro**
- IKE policy – **winRadiusPol**
- Source IP address – **192.1.1.0/24**

- Source port – standard RADIUS port **500**
- Destination IP address – **192.2.2.0/24**
- Destination port – standard RADIUS port **500**
- Encapsulation type – **transport**
- Lifetime time – **5** minutes
- Lifetime bandwidth – **100000** kilobytes
- The transmission protocol – **udp**
- Encryption request – **enabled**

```

7100 Chassis(su-config)->crypto ike-map winRadius
7100 Chassis(su-crypto-map)->proposal winRadiusPro
7100 Chassis(su-crypto-map)->policy winRadiusPol
7100 Chassis(su-crypto-map)->src 192.1.1.0/24
7100 Chassis(su-crypto-map)->src-port 500
7100 Chassis(su-crypto-map)->dst 192.2.2.0/24
7100 Chassis(su-crypto-map)->dst-port 500
7100 Chassis(su-crypto-map)->encapsulation transport
7100 Chassis(su-crypto-map)->lifetime time 5
7100 Chassis(su-crypto-map)->lifetime bandwidth 100000
7100 Chassis(su-crypto-map)->protocol udp
7100 Chassis(su-crypto-map)->request
7100 Chassis(su-crypto-map)->exit
7100 Chassis(su-config)->

```

### IPsec

For this release an IPsec default instance is configurable. You assign the IKE map **winRadius** to the IPsec default instance within IPsec default instance configuration mode. You enable IPsec on the router in global VRF router configuration mode. For this IPsec configuration example we will also enable IPsec traps.

```

7100 Chassis(su-config)->crypto ipsec default
7100 Chassis(su-crypto-ipsec-default)->ike map winRadius
7100 Chassis(su-crypto-ipsec-default)->exit
7100 Chassis(su-config)->crypto ipsec trap-enable
7100 Chassis(su-config)->crypto ipsec enable
7100 Chassis(su-config)->

```

## Terms and Definitions

[Table 11-4](#) lists terms and definitions used in this IPsec configuration discussion.

**Table 11-4 IPsec Configuration Terms and Definitions**

Term	Definition
Encapsulating Security Payload (ESP)	An IPv4 and IPv6 packet header designed to provide a mix of security services including: confidentiality, data origin authentication, connectionless integrity, depending upon supported and configured SA configuration.

**Table 11-4 IPsec Configuration Terms and Definitions (continued)**

<b>Term</b>	<b>Definition</b>
Encryption	The process of transforming information, usually referred to as plaintext, using an algorithm, called a cipher, to make it unreadable to anyone except those possessing the associated key.
ESP Authenticity	An ESP feature that ensures that the owner of the packet is who he claims to be.
ESP Confidentiality	An ESP feature that ensures that information is accessible only to those authorized to have access.
ESP Integrity	An ESP feature, also referred to as data authentication, that ensures that the contents of the packet have not been tampered with.
Hash	The Secure Hash Algorithm 1 (SHA1) hash algorithm is used during phase 1 negotiation between the SA authenticating devices.
IKE Diffie-Hellman Group	A method of exchanging keys allowing two parties that have no prior knowledge of each other to jointly establish a shared secret key over an insecure communications channel.
IKE Map	A bundling of all algorithms and parameters that make up the SA.
IKE Policy	A combination of security parameters that exist both locally and on the peer, to be used during the IKE SA negotiation.
IKE Proposal	A set of parameters applied to both Phase I and Phase II IPsec negotiations during which the two peers establish a secure connection by which they then negotiate the Phase 2 parameters.
Initial Contact	A feature that when enabled sends an initial contact message to the peer upon reboot instructing the peer to delete old SAs.
Internet Key Exchange protocol (IKE)	The protocol used to set up a Security Association (SA) in the IPsec protocol suite.
IPsec	The Internet Protocol Security Architecture, defined in RFC 4301, that provides a set of security services for traffic at the IP layer in both IPv4 and IPv6 environments.
Security Association (SA)	The establishment of shared security attributes between two network entities to support secure communication within the IPsec protocol suite.
Security Parameter Index (SPI)	An index to the security association database that helps in differentiating between two traffic streams where different encryption rules and algorithms may be in use.



## Tracked Object Manager Configuration

This document provides the following information about configuring the Tracked Object Manager on the Extreme Networks 7100-Series platform.

For information about...	Refer to page...
<a href="#">Using Tracked Object Manager in Your Network</a>	12-1
<a href="#">State Probe Configuration</a>	12-3
<a href="#">Timing Probe Configuration</a>	12-10
<a href="#">Tracked Object Configuration</a>	12-12
<a href="#">Terms and Definitions</a>	12-15

### Using Tracked Object Manager in Your Network

The Tracked Object Manager provides the ability to track local and remote objects by means of tracked objects and probes. Tracked objects monitor the state of local entities, such as interfaces. Probes monitor the state of remote entries, such as host servers. Each tracked object or probe is a container governed by a set of rules and configurable attributes used to determine the state of a collection of one or more of their respective entities. The run-time states of the tracked objects and probe are either up or down. Other states display a configuration or initiation state.

The Tracked Object Manager provides its services to client applications. An application must register with the Tracked Object Manager in order to use its monitoring services. When a client application wants to follow the state of a local service, the application registers with a tracked object. If the client application wishes to determine the state of a remote server, it creates a probe session. The Tracked Object Manager is responsible for monitoring the state of the configured objects and reports back to the registered client applications when the state of these objects change.

### Tracked Objects

Tracked objects monitor the state of different types of local entities. Currently, the **port-group** tracked object type is supported. The tracked object type dictates the rules the tracked objects obey in determining their own state. The tracked object contains the configuration of the local entities and the attributes that allow for rules which affect the state of the tracked object. Client applications register for tracked object state events and perform actions upon receiving state events for that tracked object from the Tracked Object Manager.

The port-group type of tracked object allows users to form port-groups. For example, the port-group tracked object is used by the Link-State application to monitor the state of an upstream port-group and execute an action depending on the state events received from the Tracked Object Manager for that tracked object.



## Probes

**State probes** track the availability of a remote service by actively transmitting network packets to a specified remote host. Tracked Object Manager supports three probe protocol types:

- An ICMP probe that monitors a device, by sending an ICMP ping to the IP address the probe is assigned to.
- A UDP probe that is capable of port service verification, by sending the port a UDP packet and waiting for an ICMP “Port Unreachable” response if the port is down. A UDP probe can also be configured for Application Content Verification (ACV) if the remote server supports a protocol that responds to a UDP packet, such as the UDP Echo protocol.
- A TCP probe that is capable of port service verification, by monitoring the appropriate port for services such as HTTP, Telnet, SMTP, and FTP. A TCP probe can also be configured for ACV for the verification of a layer 7 (OSI model) application running on the server.

The rules and attributes defined by the probe dictate how and when to transmit a packet to a remote host or peer. Unlike tracked objects, a probe does not contain the configuration of the remote entity. Client applications do not register for probes — they create probe sessions. The application provides the IP address and port tuple when creating the session. A probe session, like a tracked object, informs client applications of state events. It is the responsibility of the application to take an action based on the event.

**Timing probes** gather packet timing measurement for protocol packets. This is a more specialized type of probe, and it does not provide state events. Instead, each request for a probe session provides the client application with the packet transmit and receive times. Currently, ICMP and UDP timing packets are supported. The ICMP timing probes utilize the ICMP echo/reply paradigm. UDP timing probes may be configured to use the UDP-echo or DNS protocol.

The ICMP and UDP probes have the ability to gather packet timing measurements instead of monitoring the state of remote entity. Thus, each request for a probe session provides the client application with the packet transmit and receive times.

Examples of client applications that use probes include:

- Policy Based Routing (PBR), which uses an ICMP probe to monitor a next hop IP address.
- The Virtual Router Redundancy Protocol (VRRP), which uses an ICMP probe to monitor a critical IP interface.
- The IP service level agreement (SLA) application, which uses an ICMP or UDP timing probe session to collect packet timing information. This type of data collection provides the IP SLA application with the statistics it needs to perform its calculations.



**Note:** Prior to the 7100-Series Firmware Release 7.21, the tracked objects functionality was performed in policy based routing by the route map pinger feature. Both route map pinger and the previous application based fail detection have been removed from the 7100-Series firmware and have been replaced by the Tracked Object Manager feature.

## Scheduling

### Probe Session Scheduling

Probe sessions are scheduled by the Tracked Object Manager's scheduler. The scheduler rate limits the number of sessions that run every second using the leaky bucket paradigm. The scheduler's current rate limit is 200 sessions per second.

The probe attributes control the operational mechanics of the scheduler. Each probe session takes several actions during the lifetime of a particular transaction. A probe session retrieving or

releasing a system resource may cause a variation in its scheduling and it may also reduce the scheduler's rate limit for that second interval.

In addition, the scheduler's rate limiter dampens the Tracked Object Manager's CPU utilization during periods of heavy load, with the possible effect of delaying any number of the session's transactions.

Collecting packet timing information requires a constant transmit rate for the ICMP echo requests. The Tracked Object Manager supports this, but enforces a restriction on the receive wait value for ICMP timing probes. The receive wait value CANNOT be larger than the transmit interval. You configure the transmit interval and receive wait values in milliseconds for the timing probe, but the regular ICMP, UDP, and TCP probe attributes are configured in seconds.

## Tracked Object Scheduling

Tracked objects utilize the Tracked Object Manager's scheduler when the delay up or down attribute is non-zero. The scheduler acts as a countdown timer for informing the client applications of a tracked object's state change. When a tracked object's state changes, the Tracked Object Manager puts the entry on the scheduling queue for the time period dictated by the delay attribute (either up or down). If the tracked object's state reverts back to its previous state prior to the countdown timer expiring, the Tracked Object Manager does not inform the client applications of the intermittent state change.

## State Probe Configuration

**State probes** monitor the state of a remote service by actively transmitting network packets to a specified remote host. To configure a state probe:

- Create the probe by specifying a probe name and type
- Optionally configure a description to be associated with this probe
- Optionally configure a domain, host or IPv6 address DNS query type to send with this state probe
- Optionally configure an IP address or domain name to verify the DNS query response
- Optionally modify the number of consecutive failed faildetect probes that will determine when the service is declared down
- Optionally modify the interval between faildetect probes
- Optionally modify the number of successful pass detection probes that will determine when a service marked as down will be declared up
- Optionally modify the interval between pass detection probes
- Optionally specify the ACV or DNS Layer 5 protocol to use with this state probe
- Optionally modify the length of time the Tracked Object Manager will wait for a response from the monitored service before declaring that a probe request failed
- For a TCP probe, optionally modify the open interval that sets how long the Tracked Object Manager should wait for the completion of the TCP 3-way handshake
- When configuring ACV on a TCP or UDP probe:
  - Set the request string that will initiate the ACV session on the server
  - Set the reply string that will validate the server response to the request string
  - If required by the protocol being monitored, configure a close string to close the session

- Enable the probe by placing it inservice

The three state probe protocols supported by the Tracked Object Manager are ICMP, UDP, and TCP. Probe parameters are configured in probe configuration mode. You enter probe configuration mode by creating the probe in global configuration mode, specifying the name of the probe and the probe protocol. If the specified probe already exists, Tracked Object Manager enters configuration command mode for the named probe.

The state probe protocol used determines the fail detection method(s) that are available for monitoring the remote service. The fail detection methods supported for monitoring a remote service are:

- Ping
- Port Service Verification
- Application Content Verification (ACV)

Probes that do not yet exist can be assigned to monitor a service, but fail detection will not occur until the probe is created.

## Probe Parameters

Probe parameters are configurable by entering probe configuration mode from the global configuration mode.

### Description

A probe description of up to 127 printable characters can be configured. If a space character is entered, the description must be enclosed by double quotes (""). Probe descriptions display in the detailed version of the **show probe** command output.

### Application Content Verification Parameters

The following content verification parameters can be set:

- **Request String** – A string used by ACV that the Tracked Object Manager sends to the remote server to initiate verification of an application.
- **Reply String** – A string used by the Tracked Object Manager to validate the server response to the ACV request string.
- **Close String** – A string used by ACV to close a session when required by the protocol.
- **Search-Depth** – The number of characters into the server response to search for the ACV reply string. The reply string must match entirely within the search-depth.

### Fail Detection Parameters

The Tracked Object Manager uses fail detection to determine when a service that is currently declared up should be declared down. Fail detection parameters set:

- The number of consecutive failed probe attempts before Tracked Object Manager declares a remote service down
- The delay, in seconds, between probes to a remote service that is currently declared up

### Pass Detection Parameters

The Tracked Object Manager uses pass detection to determine when a service that is down should be declared up. Pass detection parameters set:

- The number of consecutive successful probes to a service currently declared down before the Tracked Object Manager declares the service up
- The delay, in seconds, between probes to a service that Tracked Object Manager currently declares down

### Common Pass/Fail Parameters

- The time, in seconds, the Tracked Object Manager waits for a response from the monitored service. If a response is received within that time, the attempt passed. If a response is not received within that time, the attempt failed.
- The time, in seconds, the Tracked Object Manager waits for the TCP 3-way handshake to complete. If the handshake completes within that time, the attempt passed. If the handshake does not complete within that time, the attempt failed.

## Fail Detection Methods

The fail detection method used determines whether the probe verifies a service, port, or application. The local application determines which fail detection methods are supported.

### Ping

A remote service can be configured for the ping failure detection method by setting the probe protocol to ICMP. The ping failure detection method can be used by all 7100-Series applications supported by the Tracked Object Manager.

## Application Content Verification

Application Content Verification (ACV) can be enabled on a port to verify the content of an application on one or more servers. ACV is a method of ensuring that the server is responding with the appropriate response given some known good request. By its nature, ACV is protocol-independent and is designed to work with any type of server that communicates via formatted ASCII text messages, including HTTP, FTP, and SMTP.

ACV can be configured on both TCP and UDP probes.

### ACV Configured On a UDP Probe

UDP is a connectionless protocol. The UDP server must have a protocol capable of responding to a UDP ACV probe request, such as the UDP Echo protocol. In the case of the UDP Echo protocol, the response is an echo of the probe request sent to the server. In this case, the configured string of the expected reply from the server is the same as the configured request string.

### ACV Configured On a TCP Probe

ACV works by sending a request to your application server and searching the response for a certain string. If it finds the string, the server is marked as Up. If the string is not found, the server is marked as Down.

For ACV verification of a TCP server application, you specify the following:

- A string that the router sends to the server. The string can be a simple HTTP command to get a specific HTML page, or it can be a command to execute a user-defined CGI script that tests the operation of the application.
- The reply that the application on each server sends back is used by the router to validate the content. In the case where a specific HTML page is retrieved, the reply can be a string that appears on the page, such as "OK". If a CGI script is executed on the server, it should return a specific response (for example, "OK") that the router can verify.

For example, if you sent the following string to your HTTP server, "HEAD / HTTP/1.1\r\nHost: www.extremenetworks.com\r\n", you could expect to get a response of a string returned similar to the following:

```
HTTP/1.1 200 OK
Date: Tue, 9 Feb 2010 20:03:40 GMT
Server: Apache/2.0.40 (Red Hat Linux)
Last-Modified: Wed, 6 Jan 2010 13:56:03 GMT
ETag: "297bc-b52-65f942c0"
Accept-Ranges: bytes
Content-Length: 2898
```

You can search for a reply string of "200 OK". This would result in a successful verification of the service.

Because ACV can search for a string in only the first 255 bytes of the response, in most HTTP cases the response will have to be in the packet's HTTP header (that is, you will not be able to search for a string contained in the web page itself).

Some protocols such as FTP or SMTP require users to issue a command to close the session after making the request. An ACV close string can be configured and sent by the Tracked Object Manager to the server to close the session.

## Preset Default ICMP Probes

Tracked Object Manager allocates a probe entry for each client application. A preset default ICMP probe for each supported application exists when you boot your system, although not all client applications may use their default probe. You cannot modify or delete default ICMP probes.

How a default ICMP probe is handled depends upon the application the default probe is associated with. Default ICMP probes associated with non-server-based applications such as policy based routing and VRRP are manually applied.

Use the **show probe default** command to display a list of the current default probes. Use the **show probe probe-name detail** command to display the description and attributes of the default probes.

## Manually Applied Default ICMP Probes

Manually applied default ICMP probes are treated the same as an administratively created ICMP probe and are provided for your convenience, should the preset parameter values meet your needs.

The Policy Based Routing (PBR) default ICMP probe must be manually applied. Use the **route-map probe** command in global configuration mode to apply the PBR default ICMP probe (**\$pbr\_default**) to monitor the specified next hop IP address. When configuring a default ICMP probe, the probe cannot be specified by name. Use the **default** keyword when configuring the default route-map probe.

The following example configures the default **\$pbr\_default** ICMP probe to monitor IP address **125.50.25.1**:

```
7100 Chassis(su-config)->route-map probe 125.50.25.1 probe-name default
```

The VRRP default ICMP probe is used to monitor remote critical IP addresses. When configuring a default ICMP probe, the probe cannot be specified by name. The VRRP default probe is configured when the **remote** keyword is specified. Use the **vrrp critical-ip** command in interface configuration mode, specifying the **remote** keyword, to apply the VRRP default probe to a critical IP interface.

This example sets the internet facing IP address **20.20.20.2** on VLAN **20** as the critical-IP address for VRRP instance **1**, sets the decrement operational priority to **100** should the interface go down, and assigns the VRRP default probe **\$vrrp\_default** to monitor the interface:

```
7100 Chassis(rw)->configure
7100 Chassis(rw-config)->interface vlan 20
7100 Chassis(rw-config-intf-vlan.0.20)->vrrp critical-ip 1 20.20.20.2 100 remote
probe-name $vrrp_default
7100 Chassis(rw-config-intf-vlan.0.20)->no shutdown
7100 Chassis(rw-config-intf-vlan.0.20)->
```

## Configuring a Probe for Policy Based Routing

The route-map manager supports the assigning of an ICMP probe to monitor a next hop IP address. The route-map facility uses the Tracked Object Manager to monitor the IP address, but the ICMP probe is not assigned to a specific route-map. If a next hop IP address is declared down, it is removed from the next hop selection process for all route-maps specifying this address as a next hop, until it is declared up again. The assigned ICMP probe will ping port 0 of the specified IPv4 or IPv6 address.

Use the **route-map probe** command in router configuration mode to assign an ICMP probe to monitor the specified next hop IP address. Create a probe, using the **probe** command. A default ICMP probe can not be specified by name. Use the **default** keyword to assign the default policy based routing ICMP probe.

This example shows how to create the ICMP probe **ICMP-PBR** and assign it to a route-map probe to monitor next hop IP addresses **101.10.1.252** and **2000::1301:0:21f:45ff:fe4d:8722**. The fail detection count is set to **5** attempts, and the fail detection interval is set to **5** seconds. The two assigned sessions are displayed:

```
7100 Chassis(su-config)->probe ICMP-PBR icmp
7100 Chassis(su-config-probe)->faildetect count 5 interval 5
7100 Chassis(su-config-probe)->inservice
7100 Chassis(su-config-probe)->exit
7100 Chassis(su-config)->route-map probe 101.10.1.252 probe-name ICMP-PBR
7100 Chassis(su-config)->route-map probe 2000::1301:0:21f:45ff:fe4d:8722
probe-name ICMP-PBR
7100 Chassis(su-config)->show probe sessions
```

```
Client Codes: P-policy based routing, S-SLB, V-VRRP, W-TWCB
               T-tracked object probe
```

...

```
Probe: ICMP-PBR, icmp
```

IP Address	Port	Status	StChngs	Last Change	Clients
101.10.1.252	0	Up	1	0h0m30s	P
2000::1301:0:21f:45ff:fe4d:8722	0	Up	1	0h0m40s	P

```
Displayed 2 sessions
```

...

```
7100 Chassis(su-config)->
```

## Configuring a Probe for VRRP

VRRP supports the assigning of an ICMP probe to monitor a remote VRRP critical IP address. If an administratively configured probe name is not specified when configuring a remote critical IP address, the default VRRP ICMP probe, `$vrrp_default` is auto-configured to monitor the remote critical IP address. See “[Preset Default ICMP Probes](#)” on page 12-6 for default ICMP probe details.

This example:

- Creates the **ICMP-VRRP** ICMP probe
- Sets the fail detection and pass detection intervals to **5** seconds
- Sets the internet facing IP address **20.20.20.2** on **VLAN 20** as the critical-IP address for VRRP instance **1**
- Sets the decrement operational priority to **10** should the interface go down
- Assigns ICMP probe **ICMP-VRRP** to monitor the interface
- Enables the interface

```
7100 Chassis(su-config)->probe ICMP-VRRP icmp
7100 Chassis(su-config-probe)->faildetect interval 5
7100 Chassis(su-config-probe)->passdetect interval 5
7100 Chassis(su-config-probe)->inservice
7100 Chassis(su-config-probe)->exit
7100 Chassis(rw)->configure
7100 Chassis(rw-config)->interface vlan 20
7100 Chassis(rw-config-intf-vlan.0.20)->vrrp critical-ip 1 20.20.20.2 10 remote
probe-name ICMP-VRRP
7100 Chassis(rw-config-intf-vlan.0.20)->no shutdown
7100 Chassis(rw-config-intf-vlan.0.20)->
```

## Configuring State Probes

This section provides details for state probe configuration on 7100-Series products.

[Table 12-1](#) lists state probe default values.

**Table 12-1 Default Tracked Object Manager Parameters**

Parameter	Description	Default Value
probe faildetect count	The consecutive number of failed attempts before the service is declared down.	3 probes
probe faildetect interval	The delay in seconds between probes to a service that is up.	10 seconds
probe passdetect count	The consecutive number of successful probes to a service marked as down before the service is declared up.	3 probes
probe passdetect interval	The delay between probes to a service marked as down.	300 seconds
probe state	The service state of a configured probe.	not-in-service

**Table 12-1 Default Tracked Object Manager Parameters (continued)**

Parameter	Description	Default Value
receive interval	The time, in seconds, the Tracked Object Manager waits for a response from the monitored service before declaring a failed probe.	10 seconds
search depth	The number of characters into the server response to search for the ACV reply string.	255 characters
SLB faildetect probe one and two	Default probe for server load balancing faildetect probe one and two.	probe one: \$slb_default probe two: empty
SLB faildetect type	The default probe behavior for this real server configuration.	probe; fail detection is active
TCP 3-way handshake interval	The interval, in seconds, the track object manager waits for the 3-way handshake to complete.	5 seconds

[Procedure 12-1](#) describes how to configure state probes. Refer to the “Tracked Object Manager Commands” chapter in the 7100-Series *CLI Reference* for more information about these commands.

#### Procedure 12-1 State Probe Configuration

Step	Task	Command(s)
1.	Create a probe, specifying the probe name and protocol type.	<b>probe</b> <i>probe-name</i> { <b>icmp</b>   <b>tcp</b>   <b>udp</b> }
2.	Optionally, configure a description to be associated with the probe.	<b>description</b> <i>description-text</i>
3.	Optionally, modify the number of consecutive failed faildetect probes that determine when the service is declared down.	<b>faildetect count</b> <i>count</i>
4.	Optionally, modify the interval between fail detection probes.	<b>faildetect interval</b> <i>seconds</i>
5.	Optionally, modify the number of successful pass detection probes that determine when a service marked as down will be declared up.	<b>passdetect count</b> <i>count</i>
6.	Optionally, modify the interval between pass detection probes.	<b>passdetect interval</b> <i>seconds</i>
7.	Optionally, specify the length of time the Tracked Object Manager waits for a response from the monitored service before declaring that a probe has failed.	<b>receive</b> <i>wait-interval</i>
8.	For a TCP probe, optionally modify the interval that sets how long the Tracked Object Manager waits for the completion of the TCP 3-way handshake.	<b>open</b> <i>wait-interval</i>
9.	When configuring ACV on a TCP or UDP probe, set the request string that will initiate the ACV session on the server.	<b>acv request</b> <i>request-string</i>



**Procedure 12-1 State Probe Configuration (continued)**

Step	Task	Command(s)
10.	When configuring ACV on a TCP or UDP probe, set the ACV validation reply string the server responds to the request string with.	<b>acv reply</b> <i>reply-string</i>
11.	When configuring ACV on a TCP probe, if required by the monitored protocol, configure a close string to close the session.	<b>acv close</b> <i>close-string</i>
12.	In probe configuration mode, optionally specify a DNS query type	<b>dns-query type</b> {[ <b>domain</b> [ <b>ip ip-address</b>   <b>name</b> ]   <b>host name</b>   <b>ipv6 ipv6-address</b> ]}
13.	In probe configuration mode, optionally specify a domain name or IP address that will be used to verify the DNS query response.	<b>dns-verify match</b> { <b>address ip-address</b>   <b>domain name</b> }
14.	In probe configuration mode, optionally specify a Layer 5 protocol to use with this probe.	<b>l5-type</b> { <b>acv</b>   <b>dns</b> }
15.	Enable the probe by placing it in service.	<b>in-service</b>

## Timing Probe Configuration

**Timing probes** gather packet timing measurements using protocol packets. The timing probe protocols currently supported by the Tracked Object Manager are ICMP and UDP.

Probe parameters are configured in probe configuration mode. You enter probe configuration mode by creating the probe in global configuration mode, specifying the name of the probe and the probe protocol and type. If the specified probe already exists, Tracked Object Manager enters configuration command mode for the named probe.

### Timing Probe Parameters

Timing probe parameters are configured in probe configuration mode using the following commands:

**Table 12-2 Configuring Timing Probe Parameters**

Command	Parameter Description	Default Value
<b>description</b>	A probe description of up to 127 printable characters can be configured. If a space character is entered, the description must be enclosed by double quotes (""). Probe descriptions are displayed by the <b>show probe probe-name detail</b> command.	None
<b>dns-query type</b>	This command allows you to specify a domain name, host name, or IPv6 address as the DNS query type to be sent with the probe.	None
<b>l5-type</b>	This command allows you to specify that the ACV or DNS Layer 5 protocol will be used with the probe.	None

**Table 12-2 Configuring Timing Probe Parameters (continued)**

Command	Parameter Description	Default Value
<b>interval</b>	This command sets the transmit rate of ICMP echo requests. The transmit rate interval must be larger than the length of time specified to wait for an ICMP echo reply (the receive wait time).  The system will not allow you to set a transmit interval less than or equal to the receive wait time, so if necessary, change the receive parameter first, before changing the interval value.	2000 milliseconds
<b>packet-options</b>	This command allows you to set the IP type of service or VLAN priority code point value to be included in the ICMP echo requests sent.	Both ToS and PCP are set to 0
<b>receive</b>	This command sets the length of time to wait for an ICMP echo reply. This receive wait time must be smaller than the transmit rate interval.  The system will not allow you to set a receive wait time that is greater than or equal to the transmit rate interval, so if necessary, change the interval parameter before changing the receive wait time.	1000 milliseconds
<b>inservice</b>	This command places the probe in service. A scheduled IP SLA entry requires the assigned timing probe to be in service; otherwise, the IP SLA tests do not return data.	out of service

## Configuring a Timing Probe for IP SLA

IP SLA requires an ICMP timing probe to test the connection to a destination IP address. You can assign the default ICMP timing probe, **\$ipsla\_default**, which uses the default ICMP timing probe settings, or you can create and assign a custom ICMP timing probe.

This example:

- Creates the **ICMP-IP-SLA** ICMP timing probe
- Sets the receive wait time to **500** milliseconds
- Sets the transmit interval to **1000** milliseconds
- Puts the ICMP timing probe into service
- Sets **ICMP-IP-SLA** as the ICMP timing probe for IP SLA tests on the destination IP address **125.50.25.1**

```
7100 Chassis(su-config)->probe ICMP-IP-SLA icmp timing
7100 Chassis(su-config-probe)->receive 500
7100 Chassis(su-config-probe)->interval 1000
7100 Chassis(su-config-probe)->inservice
7100 Chassis(su-config-probe)->exit
7100 Chassis(su-config)->sla entry 1 echo
S Chassis(su-config-sla)->destination 125.50.25.1 probe ICMP-IP-SLA
```

To run the IP SLA entry, you must schedule it. For information about scheduling the IP SLA entry, see [Chapter 13, IP SLA Configuration](#).

## Procedure

[Procedure 12-2](#) lists the steps to configure an ICMP or UDP timing probe. Refer to the “Tracked Object Manager Commands” chapter in the S-Series CLI Reference for details about using these commands.

### Procedure 12-2 Timing Probe Configuration

Step	Task	Command(s)
1.	In global configuration mode, create the timing probe and enter probe timing configuration mode.	<code>probe probe-name {icmp   udp} timing</code>
2.	In probe timing configuration mode, optionally configure a description of the probe.	<code>description string</code>
3.	In probe timing configuration mode, optionally change the transmit interval from the default of 2000 milliseconds.	<code>interval millisecs</code>
4.	In probe timing configuration mode, optionally change the receive wait time from the default of 1000 milliseconds.	<code>receive millisecs</code>
5.	In probe timing configuration mode, optionally set the IP type of service or VLAN priority code point value to be included in the ICMP echo requests sent.	<code>packet-options {ip-tos tos   vlan-pcp pcp}</code>
6.	In probe timing configuration mode, optionally specify a DNS query type	<code>dns-query type {[domain [ip ip-address   name]   host name   ipv6 ipv6-address]}</code>
7.	In probe configuration mode, optionally specify a Layer 5 protocol to use with this probe.	<code>l5-type {acv   dns}</code>
8.	In probe timing configuration mode, put the probe inservice.	<code>inservice</code>
9.	Display information about the probe in any command mode.	<code>show probe [probe-name [detail   session]]</code>  <code>show probe sessions</code>  <code>show probe default</code>

## Tracked Object Configuration

Tracked objects monitor the state of different types of local entities. Currently, the **port-group** tracked object type is supported. The port-group tracked object allows you to monitor the line protocol status of a group of ports.

Tracked object parameters are configured in tracked object configuration mode. You enter tracked object configuration mode by creating the tracked object in global configuration mode, specifying the name of the tracked object and the tracked object type. If the specified tracked object already exists, Tracked Object Manager enters configuration command mode for the named object.

### Tracked Object Parameters

Tracked object parameters are configured in tracked object configuration mode using the following commands:

**Table 12-3 Configuring Tracked Object Parameters**

Command	Parameter Description	Default Value
<b>port</b>	Valid for port-group tracked objects only. Specifies ports to be added to the port group for tracking.	None
<b>threshold count</b>	Valid for port-group tracked objects only. This command sets the port group threshold counts which control the up and down state of the port group tracked object.	
<b>up</b>	The tracked object changes to the “up” state if the number of “up” ports is greater than or equal to the <b>up</b> count value. Up count value can range from 1 to 255.	Up: 1
<b>down</b>	The tracked object changes to the “down” state when the number of “up” ports is less than or equal to the <b>down</b> count value. Down count value can range from 0 to 254.	Down: 0
	The down count must be smaller than the up count. The up count must be greater than the down count.	
<b>delay up   down</b>	This command configures the amount of time for the Tracked Object Manager to wait prior to informing client applications of a state change, either “up” or “down”. Value can range from 1 to 180 seconds.	Up: 3 seconds Down: 3 seconds
<b>description</b>	A description of up to 127 printable characters can be configured. If a space character is entered, the description must be enclosed by double quotes (“”).  Tracked object descriptions are displayed by the <b>show track track-name detail</b> command.	None
<b>inservice</b>	This command enables the tracked object .	Disabled

## Procedure

[Procedure 12-3](#) on page 12-13 lists the steps to configure a tracked object. Refer to the “Tracked Object Manager Commands” chapter in the S-Series CLI Reference for details about using these commands.

### Procedure 12-3 Port Group Tracked Object Configuration

Step	Task	Command(s)
1.	In global configuration mode, create a tracked object, specifying the type, and enter tracked object configuration mode.	<b>track track-name port-group</b>
2.	In port-group tracked object configuration mode, configure the ports to be included in the port-group object.	<b>port port-string</b>
3.	Optionally, in tracked object configuration mode, change the default delay values.	<b>delay { [up secs]   [down secs] }</b>
4.	Optionally, in port-group tracked object configuration mode, change the default threshold values.	<b>threshold count { [up count]   [down count] }</b>
5.	Optionally, in tracked object configuration mode, specify a description for the tracked object.	<b>description string</b>

**Procedure 12-3 Port Group Tracked Object Configuration (continued)**

Step	Task	Command(s)
6.	In tracked object configuration mode, enable the tracked object.	<code>inservice</code>
7.	Display information about the tracked object in any command mode.	<code>show track [track-name [detail]]</code>

**Example**

This example creates a port-group tracked object named `ls_group`, changes the defaults for up/down status message delays for the object, gives it a description, puts it in service, then configures the Link-State application to use that tracked object and associate it with downstream ports.

```

7100 Chassis(su)->configure
7100 Chassis(su-config)->track ls_group port-group
7100 Chassis(su-config-track-obj)->delay up 5 down 5
7100 Chassis(su-config-track-obj)->description "link-state group1"
7100 Chassis(su-config-track-obj)->port tg.2.1-4
7100 Chassis(su-config-track-obj)->inservice
7100 Chassis(su-config-track-obj)->exit
7100 Chassis(su-config)->exit
7100 Chassis(su)->show track ls_group detail

Description: link-state group1
Track ls_group
  Port-Group tg.2.1-4
    4 ports used [0 up, 4 down]
    threshold up 1, down 0
    speed 0 [aggregate]
  Status is Down
    2 changes, last change 2d18h44m01s
  Delay up 5 seconds, down 5 seconds
  Registrants:
    Link-state

Displayed 1 tracked objects
7100 Chassis(su)->set link-state track ls_group downstream ge.1.1-5
7100 Chassis(su)>show link-state ls_group detail

Link-state ls_group
  Ports
    Uplinks: tg.2.1-4
    Downlinks: ge.1.1-5
  State is Down, Last action shutdown downlinks
    3 state changes, last change 2d18h47m21s ago

Displayed 1 link-state entries

```

## Terms and Definitions

Table 12-4 lists terms and definitions used in this Tracked Object Manager configuration discussion.

**Table 12-4 Tracked Object Manager Terms and Definitions**

Term	Definition
Tracked Object Manager	The Tracked Object Manager provides the ability to track local and remote objects by means of tracked objects and probes.
state probe	A probe of protocol type ICMP, UDP, or TCP that tracks the availability of a remote service, by actively transmitting network packets to a specified remote host.
timing probe	A probe of protocol type ICMP that gathers packet timing measurements for protocol packets. Timing probes do not provide state events. Instead, each request for a probe session provides the client application with the packet transmit and receive times.
tracked object	Tracked objects monitor the state of different types of local entities.
client application	An application that uses the objects provided by Tracked Object Manager
probe session	An entity consisting of an IP/port tuple created by a client application using a particular probe.
<b>State Probe Terms</b>	
server port verification	A state probe fail detection method used by server load balancing and TWCB to assure that the remote server is up.
application content verification (ACV)	A state probe fail detection method for the verification of application content on a server.
ICMP ping	A fail detection method that sends a ping packet to the IP address of the remote service.
default ICMP probe	A preset probe configured for each of the supported local applications.
DNS query type	A domain or host name or IPv6 address that is sent with the state probe.
DNS verify type	A domain name or IP address that is used to verify the DNS query response.
close string	A string used by ACV to close a session.
reply string	A string used by the Tracked Object Manager to validate the server response to the ACV request string.
request string	A string used by ACV that the local application sends to the remote server to initiate verification of an application.
search depth	A numeric value that specifies the number of characters to search within an ACV response for the ACV reply string.
faildetect count	The number of consecutive failed probe attempts before Tracked Object Manager declares a remote service down.
faildetect interval	The delay, in seconds, between probes to a remote service that is currently declared up.
probe one and two	Up to two probes, that can be a default probe or administratively created probe, labelled <b>one</b> and <b>two</b> , applied to a server context.
fail detection type	Specifies whether or not fail detection is active in the current server context.
L5-type	Specifies a ACV or DNS Layer 5 protocol to use with this state probe.

**Table 12-4 Tracked Object Manager Terms and Definitions (continued)**

<b>Term</b>	<b>Definition</b>
open interval	The time, in seconds, the Tracked Object Manager waits for the TCP 3-way handshake to complete.
passdetect count	The number of consecutive successful probe attempts to a service currently declared down before the Tracked Object Manager declares the service up.
passdetect interval	The delay, in seconds, between probes to a remote service that is currently declared down.
receive interval	The time, in seconds, the Tracked Object Manager waits for a response from the remote service before declaring a failed probe.
<b>Timing Probe Terms</b>	
DNS query type	A domain or host name or IPv6 address that is sent with the timing probe.
L5-type	Specifies a ACV or DNS Layer 5 protocol to use with this timing probe.
transmit interval	The transmit rate of ICMP echo requests. The transmit rate interval must be larger than the length of time specified to wait for an ICMP echo reply (the receive wait time).
receive wait time	The length of time to wait for an ICMP echo reply. The receive wait time must be smaller than the transmit rate interval.
<b>Tracked Object Terms</b>	
threshold counts	The port-group threshold counts which control the “up” and “down” state of the port group tracked object. Specifies the number of ports that must be in an up state for the port-group object to be considered up, and the number of ports that must be in a down state for the port-group object to be considered down.
delay	The amount of time for the Tracked Object Manager to wait prior to informing client applications of a state change, either “up” or “down”.

## IP SLA Configuration

This document provides the following information about configuring IP SLA entries on the Extreme Networks 7100-Series platform.

For information about...	Refer to page...
<a href="#">Using IP SLA in Your Network</a>	13-1
<a href="#">Configuring IP SLA</a>	13-4

### Using IP SLA in Your Network

Service level agreements (SLAs) between Enterprise IT departments and end-users provide service guarantees for business critical applications. These agreements require performance monitoring of the network on a continual basis. The IP SLA feature allows you to configure, schedule, and monitor end-to-end packet timing measurements.

Use these timing measurements to understand how each service is performing on your network and, if necessary, deploy additional network applications more effectively or troubleshoot existing applications.

IP SLA collects, aggregates, and provides the ability to store the statistics gathered from the timing measurements for each session request. IP SLA performs the session requests through the Tracked Object Manager's ICMP timing probe feature. The Tracked Object Manager transmits ICMP echo packets to the destination provided by IP SLA and reports the timing information back to IP SLA.

IP SLA uses the timing information to calculate round-trip delay. Additional information provided by the Tracked Object Manager indicates if a packet was lost or is out of order. IP SLA provides a small storage area to keep the timing information for statistical modeling of the network.

### Constraints and Limitations

IP SLA uses the Tracked Object Manager's ICMP timing probe functionality to perform ICMP echo requests to capture timing information. The Tracked Object Manager throttles the amount of traffic it generates (128 ICMP echo requests per second, 64 UDP or TCP requests per second) to limit the amount of CPU used. Because the Tracked Object Manager is a shared resource, the throttle applies to all applications that use the Tracked Object Manager. The Tracked Object Manager further restricts statistic gathering probes to one request every tenth of a second (maximum of ten requests per seconds).

IP SLA allocates memory to support eight entries. The default functionality of the entry consumes one session from the Tracked Object Manager's resource pool during a scheduled test. The Tracked Object Manager has a limit of 2000 sessions. If you configure an IP SLA entry to monitor



paths, the number of sessions required for the entry increases by the product of the number of paths multiplied by the number of hops.

IP SLA also allocates 30,000 statistical entries, with each entry using approximately 240 bytes, for the distribution and history mechanisms. The statistical entry contains the data for the round-trip-time metric. The number of statistical entry resources required to start the test is dependent on the distribution count, number of paths and hops, and the number of history buckets. The calculation is as follows:

$$(\text{distribution count} + \text{history buckets}) * (\text{paths} * \text{hops} + 1)$$

For more information about the memory used by IP SLA entries and statistical entries, see [“System Resources Affected by IP SLA”](#) on page 13-3.

## Monitoring Paths

The IP SLA **monitor** command allows you to probe up to eight hops along four different equal-cost paths to the destination host. Path monitoring is restricted to the abilities of traceroute, which uses UDP packets. IP SLA uses the traceroute program to determine the hops along the different equal-cost paths. The traceroute program runs after IP SLA creates the probe session for the destination host. When there is more than one equal-cost path to the destination, the traceroute program runs for each path sequentially. After the traceroute completes for each path, IP SLA creates probe sessions for up to eight hops along that path. Each of these probe sessions stores the round-trip-time data.

## Scheduling Tests

Use the IP SLA schedule mode to start and stop tests. The schedule mode includes options to set the following:

- The start time of the test cycle
- The number of tests in the test cycle
- The duration of each test in the test cycle
- The interval between tests in the test cycle
- Whether the test cycle is repeated
- You can schedule the test to repeat and the delay between each test.

When the test starts, it starts a new statistical collection and, if configured, a history collection. In both cases, the IP SLA scheduler clears the old collection of all of its data. The scheduler then reserves all of the statistic entries it needs to perform the test. After the resources are reserved, the application creates the probe session with the Tracked Object Manager.

## Reported Statistics

IP SLA calculates the round-trip time from the timing information provided in the packet. If the packet is late or out-of-order, IP SLA updates the corresponding counter for the entry and disregards the rest of the data.

If IP SLA needs to correlate this data into a distribution or history collection, the event is stored accordingly.

If an ICMP timing probe assigned to an IP SLA entry is placed out of service while the IP SLA entry is running, the IP SLA entry will not report statistics until the ICMP timing probe is placed back in service.

The statistics for each test is stored in an IP SLA entry collection, which you can view with the **show sla** commands. For more information, see “[IP SLA Display Commands](#)” on page 13-7.

The statistics are stored locally on each blade. If a failover occurs on the blade, the IP SLA data is lost.

## Measurements

IP SLA uses the timestamps in the ICMP timing packet provided by the Tracked Object Manager to measure the round-trip delay.

- **EchoTxTime**—The timestamp added to the ICMP echo request just prior to calculating the IP checksum and transmitting the packet to its destination.
- **EchoRxTime**—The timestamp of the arrival time of the ICMP echo request at the destination. The Tracked Object Manager adds this timestamp to the ICMP echo reply prior to the ReplyTxTime timestamp.
- **ReplyTxTime**—The timestamp added to the ICMP echo reply just prior to calculating the IP checksum and transmitting the packet back to the source of the ICMP echo request.
- **ReplyRxTime**—The timestamp of the arrival time of the ICMP echo reply.

With the timing information, you can create a statistical model the network. For all measurements, the minimum, average, and maximum values are included in the output of any collection. Also included is the sum of the values for each category and the sum of squares. The sums of the squares partition the variance into manageable portions.

To calculate the standard deviation with the data provided, use the computing formula for standard deviation.

$$\text{Standard deviation} = \sqrt{(\text{sumsquared} \div \text{samples}) - ((\text{sum} \div \text{samples})^2)}$$

The following are the round-trip delay calculations.

- Round-trip delay (RTD): **ReplyRxTime – EchoTxTime**
- Wire-time RTD (WRTD): **(( ReplyRxTime – EchoTxTime) – (ReplyTxTime – EchoRxTime))**
- Host-time RTD (HRTD): **RTD – WRTD**

The calculation of the round-trip delay is always available. If the **EchoRxTime** and **ReplyTxTime** are available, the wire-time and host-time round-trip delay are calculated.

## System Resources Affected by IP SLA

IP SLA affects the following system resources:

- **Memory**—The memory consumed by IP SLA is approximately 8 Megabytes. The 30,000 statistical data entries consume a little over 7 Megabytes while the eight IP SLA entries consume less than 1 Megabyte.
- **CPU**—Through its use of the Tracked Object Manager, IP SLA may have a significant effect on CPU utilization. The number of probe sessions and frequency of ICMP echo requests sent on behalf of these sessions affect the CPU. The functioning of IP SLA and its scheduler does not have a large effect on the CPU, with the exception of displaying a large amount of data collected for an IP SLA entry.
- **Storage and persistence**—The information pertaining to the configuration is persistent and it is stored in nonvolatile storage. The timing statistics are stored in RAM and IP SLA DOES NOT distribute this data to the other line cards.

- **Network bandwidth**—If you configure path monitoring, IP SLA uses the traceroute program to determine the path to the destination host. Indirectly, IP SLA creates probe sessions with the Tracked Object Manager, which uses the network bandwidth to perform its tasks.

## IP SLA Syslog Messages

IP SLA generates the following Syslog messages:

- **Queue Overflow**—Occurs when, due to heavy usage, IP SLA cannot queue session statistics data from the Tracked Object Manager.

Logging Level 5 (warning): Queue overflow

<timestamp> <prefix> [vrf.process] <action> IP SLA application: Timing statistics queue overflow. Data dropped %d.

**Example:**

```
<166>Jan 13 16:36:20 10.21.130.55 Ipsla [1.tTrckStats] IP SLA application:
Timing statistics queue overflow. Data dropped 5.
```

- **Scheduler Started**—Occurs after the boot process and indicates IP SLA will start scheduling tests. The IP SLA scheduler does not start until the Tracked Object Manager starts.

Logging Level 7 (information): Scheduler started

<timestamp> <prefix> [vrf.process] INFO: Scheduler Started

**Example:**

```
FRI NOV 16 10:13:12 2012 IPsla[1.tIpSlaSchd] INFO: Scheduler Started
```

- **Destination Not Configured**—Occurs when the IP SLA entry attempts to start a test, but the destination IP address was not configured.

Logging Level 6 (notice): Destination not configured

<timestamp> <prefix> [vrf.process] NOTICE: Scheduled entry <n> could not run. Check destination IP address.

**Example:**

```
FRI NOV 16 10:17:46 2012 IPsla[1.tIpSlaSchd] NOTICE: Scheduled entry 4 could
not run. Check destination IP address.
```

- **Probe Deletion**—Occurs if you delete an ICMP timing probe while IP SLA entries using that probe are running test cycles.

Logging Level 6 (notice): Probe deletion

<timestamp> <prefix> [vrf.process] NOTICE: Scheduler [entry <n>]: Probe <name> deleted, ending test cycle

**Example:**

```
FRI NOV 16 10:25:09 2012 IPsla[1.tIpSlaEtsc] NOTICE: Scheduler [entry 4]: Probe
ICMP-TIMING deleted, ending test cycle
```

## Configuring IP SLA

Once you have determined how to implement IP SLA on your S-Series device, the following sections will help you configure IP SLA.

For information about...	Refer to page...
<a href="#">Default Settings</a>	13-5
<a href="#">IP SLA Configuration Procedure</a>	13-6
<a href="#">Example IP SLA Configuration</a>	13-6
<a href="#">IP SLA Display Commands</a>	13-7

## Default Settings

This section provides details for IP SLA configuration on 7100-Series devices.

[Table 13-1](#) lists IP SLA default values.

**Table 13-1 Default IP SLA Values**

Parameter	Description	Default Value
collections	The number of tests for which statistical data is kept. The default value clears the data prior to each test execution.	1 collection
distribution count	The number of distributions of measured statistics.	0 distributions
distribution interval	The length of a distribution.	25 milliseconds
history ageout	The amount of time before the application frees the resources accumulated for a history collection. A value of zero indicates the application does not free any of the resources.	0 minutes
history bucket	The number of timing information storage units in a history collection.	15 buckets
history samples	The number of samples to collect in a static-depth bucket.	16 samples
history interval	The length of a timed-depth bucket.	30 seconds
history collections	The number of history collections to maintain. The default value indicates that no storage information is kept.	0 collections
history collections wrap	Indicates whether the storage history wraps when the number of collections exceeds the value specified by <b>history collections</b> .	no
monitor path-count	The number of paths to keep track of.	0 paths
monitor hop-count	The number of hops to keep track of.	1 hop
duration	The length of a test.	30 seconds
repetitions	The number of times a test is executed.	1
recurrence	The time between test cycles.	0 seconds
frequency	The time between each test in the test cycle.	30 seconds

## IP SLA Configuration Procedure

[Procedure 13-1](#) describes how to configure an IP SLA entry. Refer to the “IP SLA” chapter in the 7100-Series *CLI Reference* for more information about these commands.

### Procedure 13-1 IP SLA Configuration

Step	Task	Command(s)
1.	In configuration mode, create an IP SLA entry, specifying the entry number and the entry type. At this time, echo is the only supported entry type.	<b>config</b> <b>sla entry</b> <1-8> <b>echo</b>
2.	Set the destination and the ICMP timing probe to be used by the IP SLA entry. You can use the default ICMP timing probe or you can create an ICMP timing probe using the Tracked Object Manager.  For more information about creating ICMP timing probes, see “ <a href="#">Timing Probe Configuration</a> ” on page 12-10.	<b>destination</b> <i>IP-address</i> <b>probe</b> { <b>default</b>   <i>probe-name</i> } [ <b>port</b> <i>port</i> ]
3.	Set the number of statistical collections for the IP SLA entry.	<b>collections</b> <1-10>
4.	Set the distribution count and interval for the IP SLA entry.	<b>distribution</b> {[ <b>count</b> <i>count</i> ] [ <b>interval</b> <i>milliseconds</i> ]}
5.	Set the history collection information for the IP SLA entry.	<b>history ageout</b> <i>minutes</i> <b>history buckets</b> <i>buckets</i> { [ <b>samples</b> <i>samples</i> ]  [ <b>interval</b> <i>seconds</i> ] } <b>history collections</b> <i>collections</i> [ <b>wrap</b> ]
6.	Set the number of hops and paths to be monitored by the IP SLA entry.	<b>monitor</b> [ <b>hop-count</b> <i>count</i> ] [ <b>path-count</b> <i>count</i> ]
7.	Exit the IP SLA entry configuration mode	exit
8.	Enter the IP SLA schedule mode.	<b>sla schedule</b>
9.	Schedule the IP SLA entry.	<b>entry</b> <i>ip-sla-entry</i> { [ <b>start</b> { { [ <b>time</b> <yyyymmdd:hh:mm:ss>]   [ <b>now</b> ]   [ <b>after</b> <5-300>}}][ <b>duration</b> <30-3600>} [ <b>frequency</b> <30-3600>} [ <b>recurrence</b> <120-7776000>} [ <b>repetitions</b> <1-10>} [ <b>reset</b> ]   [ <b>stop</b> ] }

### Example IP SLA Configuration

IP SLA requires an ICMP timing probe to test the connection to a destination IP address. You can assign the default ICMP timing probe, **\$ipsla\_default**, which uses the default ICMP timing probe settings, or you can create and assign a custom ICMP timing probe.

This example:

- Creates IP SLA entry 2
- Sets the destination IP address for the ICMP requests
- Sets the ICMP timing probe to the default ICMP timing probe, **\$ipsla\_default**
- Sets the number of collections to 2

- Sets the distribution count to 4 and the interval to 20 milliseconds
- Sets the history ageout time to 18 minutes
- Sets the history buckets to 10 buckets with 22 samples in each bucket
- Sets the history collections to 2 and enables collection wrapping
- Sets the hops and paths to monitor (8 hops, 2 paths)
- Exits the IP SLA configuration mode
- Enters the IP SLA schedule mode
- Schedules IP SLA entry 2 to start test cycles 60 seconds after the command is entered. Each test cycle includes five tests that are each 60 seconds long. The interval between the tests in the test cycle is set to 60 seconds. The next test cycle will begin one second after the current test cycle ends.

```

7100 Chassis(su-config)->sla entry 2 echo
S Chassis(su-config-sla)->destination 1.1.1.1 probe default
7100 Chassis(su-config-sla)->collections 2
7100 Chassis(su-config-sla)->distribution count 4 interval 20
7100 Chassis(su-config-sla)->history ageout 18
7100 Chassis(su-config-sla)->history buckets 10 samples 22
7100 Chassis(su-config-sla)->history collections 2 wrap
7100 Chassis(su-config-sla)->monitor hop-count 8 path-count 2
7100 Chassis(su-config-sla)->exit
7100 Chassis(su-config)->sla schedule
S Chassis(su-config-sla-sched)->entry 2 start after 60 recurrence 601 duration 60
repetitions 5 frequency 60

```

## IP SLA Display Commands

Table 13-2 lists the IP SLA show commands.

**Table 13-2 IP SLA Show Commands**

Task	Command
To display configuration and schedule information for all IP SLA entries.	<b>show sla</b>
To display the schedule information for all IP SLA entries.	<b>show sla scheduler</b>
To display configuration and schedule information for a specific IP SLA entry.	<b>show sla entry entry number</b>
To display detailed configuration and schedule information for a specific IP SLA entry.	<b>show sla entry entry number detail</b>
To display statistical distribution data for a specific IP SLA entry. Use the <b>collection</b> option to filter the output.	<b>show sla entry entry number distribution [collection collection number [destination   [path path number [hop hop number]]]]</b>
To display statistical history data for a specific IP SLA entry. Use the <b>collection</b> option to filter the output.	<b>show sla entry entry number history [collection collection number [destination]   [path path number [hop hop number [bucket bucket number]]]]</b>

**Table 13-2 IP SLA Show Commands**

Task	Command
To display statistical summary data for a specific IP SLA entry. Use the <b>collection</b> option to filter the output.	<b>show sla entry</b> <i>entry number</i> <b>summary</b> [ <b>collection</b> <i>collection number</i> [ <b>destination</b>   [ <b>path</b> <i>path number</i> [ <b>hop</b> <i>hop number</i> ]]]]]
To display the number of system resources left.	<b>show limits application sla-entry-data</b>

Refer to the *Extreme Networks 7100-Series CLI Reference* for a description of the output of each command.

## Power over Ethernet Configuration

This chapter provides information about configuring and monitoring Power over Ethernet (PoE) on the 7100-Series devices.

### Important Notice

This section applies only to PoE-equipped 7100-Series devices. Consult the *Hardware Installation Guide* shipped with your product to determine if it is PoE-equipped.

For information about...	Refer to page...
<a href="#">How to Use PoE in Your Network</a>	14-1
<a href="#">Implementing PoE</a>	14-1
<a href="#">Configuring PoE</a>	14-3

## How to Use PoE in Your Network

PoE, defined in IEEE standards 802.3af and 802.3at, refers to the ability to provide 54 Vdc (for 802.3at) or 48 Vdc (for 802.3af) operational power through an Ethernet cable from a switch or other device that can provide a PoE-compliant port connection to a powered device (PD).

Examples of PDs include:

- Voice over IP devices such as PoE-compliant digital telephones
- Devices that support Wireless Application Protocol (WAP) such as wireless access points and security cameras

Ethernet implementations employ differential signals over twisted pair cables. This requires a minimum of two twisted pairs for a single physical link. Both ends of the cable are isolated with transformers blocking any DC or common mode voltage on the signal pair. PoE exploits this fact by using two twisted pairs as the two conductors to supply a direct current to a PD. One pair carries the power supply current and the other pair provides a path for the return current.

Using PoE allows you to operate PDs in locations without local power (that is, without AC outlets). Having such a network setup can reduce the costs associated with installing electrical wiring and AC outlets to power the various devices.

## Implementing PoE

You can configure PoE on your PoE-compliant Extreme Networks device through the CLI-based procedures presented in the section “[Configuring PoE](#)” on page 14-3. As part of your plan to implement PoE in your network, you should ensure the following:

- The power requirements of your PDs are within the limits of the PoE standards.



- Your PoE-compliant Extreme Networks device can supply enough power to run your PDs. See [Table 14-1](#) for power ranges based on each device class.

**Table 14-1 PoE Powered Device Classes**

Class	Power Output at Port	Power Range Used by Device
0	15.4 watts	0.44 to 12.95 watts
1	4.0 watts	0.44 to 3.84 watts
2	7.0 watts	3.84 to 6.49 watts
3	15.4 watts	6.49 to 12.95 watts
4	Reserved (802.3af)	Treat as class 0 (802.3af)

If SNMP traps are enabled, the Extreme Networks device generates a trap to notify the network administrator if a power state occurs on a PD (for example, when a PD is powered up or unplugged)

If insufficient power is available for an attached PD, the corresponding port LED on the Extreme Networks device turns amber, when the port is in PoE mode. The LED also turns amber if a PoE fault occurs (for example, a short in the Ethernet cable).

## Allocation of PoE Power to Modules

The switch firmware determines the power available for PoE based on hardware configuration, power supply status, and power supply redundancy mode. The system calculates and reserves the correct amount of power required by the installed hardware components and then makes the balance of power available for PoE. When any change is made to the hardware configuration, power supply status, or redundancy mode, the firmware recalculates the power available for PoE.

On the 7100-Series switch, you can manually configure the maximum percentage of PoE power available to the chassis as a percentage of the total installed PoE power with the **set inlinepower available** command. If the power needed or requested exceeds the power available, the system will generate a trap to notify the system manager, if traps are enabled.

The power available for PoE is distributed based on the configured allocation mode, set with the **set inlinepower mode** command:

- **Automatic** mode, in which available power is distributed evenly to PoE-capable modules based on PoE port count. (This is the default mode.) Any change in available power, due to a change in power supply status or redundancy mode or to the addition or removal of modules, will trigger an automatic redistribution of power.
- **Manual** mode, in which the power budget for each PoE-capable module is manually configured, using either CLI commands or the MIBs. The sum of the wattage configured for each module cannot exceed the total power available on the switch for PoE.

The power budget for each PoE-capable module can be configured manually on the 7100-Series switch with the command **set inlinepower assigned**.

The configured wattage assignments are used to calculate each slot's percentage of total available power. If the total available PoE power is reduced, a redistribution of available power will occur, applying the calculated percentages.

## When Manual Mode is Configured

When manual distribution mode is configured, if a PoE module is added to the switch, the PoE power budget for existing modules will **not** be recalculated. The new module will have a power budget of zero until it is manually provisioned. Since the sum of the manually provisioned

wattages cannot exceed the total system power available, it may be necessary to adjust existing budgets to free up power for the new module.

When a PoE module is removed from a switch configured with manual power distribution mode, the PoE budget for each module will **not** be recalculated, based on the assumption that the module removed will be replaced with a new module that should receive the same amount of PoE power.

As noted above, if the total available PoE power is reduced, the power will automatically be redistributed based on applying the calculated percentages. If an additional PoE supply is installed, there is no impact on the assigned PoE since specific wattages have been assigned to each module. Only the “Total Power Detected” value will change. The extra PoE power, however, is available for further redistribution manually.

## Management of PoE Power to PDs

For each PoE-capable module or switch, you can configure how its PoE controller makes power available to attached powered devices (PDs). On a per module basis, you can configure:

- **Real-time** mode, in which the PoE controller calculates the power needed by a PD based on the actual power consumption of the attached devices.
- **Class** mode, in which the PoE controller manages power based on the IEEE 802.3af/.3at definition of the class limits advertised by the attached devices. In this mode, the maximum amount of power required by a device in the advertised class is reserved for the port, regardless of the actual amount of power being used by the device.

Power management to PDs is configured with the command **set inlinepower management**. PoE classes are defined in [Table 14-1](#) on page 14-2.

## Configuring PoE

Once you have determined how to implement PoE on your 7100-Series device, the following sections will help you configure PoE.

For information about...	Refer to page...
<a href="#">Default Settings</a>	<a href="#">14-3</a>
<a href="#">PoE Configuration Procedure</a>	<a href="#">14-4</a>
<a href="#">PoE Display Commands</a>	<a href="#">14-6</a>

## Default Settings

[Table 14-2](#) lists PoE parameters and their default values.

**Table 14-2 Default PoE Parameter Values**

Parameter	Description	Default Value
Total Power Available	The percentage of total power available that a chassis can withdraw from the total power detected.	100
Power Allocation Mode	The allocation mode for system power available for PoE.	auto

**Table 14-2 Default PoE Parameter Values (continued)**

Parameter	Description	Default Value
Power Trap Status	Whether an SNMP trap message is sent when the status of the chassis PoE power supplies or the PoE system redundancy changes.	disable
Usage Trhld	The PoE usage threshold on a module or a Standalone.	75%
PSE Trap Status	Whether an SNMP trap message is sent whenever the status of a module's ports changes, or whenever the module's PoE usage threshold is crossed.	disable
Mgmt Mode	The PoE management mode.	realtime
Admin Status	Whether PoE is enabled on the port.	auto
Priority	Which ports continue to receive power in a low power situation.	low
Power Limit	The maximum power, in milliwatts, allowed on a port.	15400 mW
Power Capability-Selection	The PoE mode selected for the port.	8023af

## PoE Configuration Procedure

[Procedure 14-1](#) describes how to configure PoE. Unspecified parameters use their default values.

### Procedure 14-1 PoE Configuration

Step	Task	Command(s)
1.	Configure PoE parameters on ports to which PDs are attached. <ul style="list-style-type: none"> <li><b>admin</b> — Enables (<b>auto</b>) or disables (<b>off</b>) PoE on a port. The default setting is <b>auto</b>.</li> <li><b>priority</b> — Sets which ports continue to receive power in a low power situation. If all ports have the same priority and the system has to cut power to the PDs, the PDs attached to the lowest numbered ports have the highest priority for receiving power. The default setting is <b>low</b>.</li> <li><b>type</b> — Associates an alias with a PD, such as "siemens phone."</li> <li><b>powerlimit</b> — Sets the maximum power, in milliwatts, allowed on a port. Valid values are 0–15400 for 802.3af and 0–34000 for 802.3at. How this parameter is set can affect the class of PD that can be attached to the port.</li> <li><b>capability</b> — Sets the PoE mode for the port to 8023af (15.4W maximum power) or 8023at (34.0W maximum power).</li> </ul>	<pre>set port inlinepower port-string {[admin {off   auto}] [priority {critical   high   low}] [type type] [powerlimit powerlimit] [capability capability]}</pre>

### Procedure 14-1 PoE Configuration (continued)

Step	Task	Command(s)
	<p>Use the <b>clear</b> command to set the port's PoE parameters back to the default settings.</p> <ul style="list-style-type: none"> <li>• <b>admin</b> — auto</li> <li>• <b>priority</b> — low</li> <li>• <b>type</b> — null</li> <li>• <b>powerlimit</b> — 15400</li> <li>• <b>capability</b> — 8023af</li> </ul>	<b>clear port inlinepower</b> <i>port-string</i> {[admin] [priority] [type] [powerlimit] [capability]}
2.	<p>(Optional) Enable an SNMP trap message to be sent when the status of the chassis PoE power supplies or the PoE system redundancy changes.</p> <p>Use the <b>clear</b> command to reset chassis power trap messaging back to the default state of disabled.</p>	<b>set inlinepower powertrap</b> {disable   enable}  <b>clear inlinepower powertrap</b>
3.	<p>(Optional) Enable an SNMP trap message to be sent whenever the status of a module's ports changes, or whenever the module's PoE usage threshold is crossed.</p> <p>Use the <b>clear</b> command to reset PoE trap messaging for a module back to default state of disabled.</p>	<b>set inlinepower psetrap</b> {disable   enable} <i>module-number</i>  <b>clear inlinepower psetrap</b> <i>module-number</i>
4.	<p>(Optional) Set the PoE usage threshold on a module. Valid values are 1–99 percent. If your 7100-Series device is a Standalone, specify 1 as the <i>module-number</i>.</p> <p>Use the <b>clear</b> command to reset the PoE usage threshold on a specified module to the default value of 75 percent.</p>	<b>set inlinepower threshold</b> <i>usage-threshold</i> <i>module-number</i>  <b>clear inlinepower threshold</b> <i>module-number</i>
5.	<p>(Optional) Set the percentage of total power available that a chassis can withdraw from the total power detected.</p> <p>Use the <b>clear</b> command to reset the percentage of the total power available to a chassis to the default value of 100.</p>	<b>set inlinepower available</b> <i>max-percentage</i>  <b>clear inlinepower available</b>
6.	<p>(Optional) Set the PoE management mode on a specified module.</p> <ul style="list-style-type: none"> <li>• <b>realtime</b> — Manages power based on the actual power consumption of the ports.</li> <li>• <b>class</b> — Manages power based on the IEEE 802.3af definition of the class upper limit for each attached PD. In this mode, the maximum amount of power required by a PD in the advertised class is reserved for the port, regardless of the actual amount of power being used by the device.</li> </ul> <p>If your 7100-Series device is an Standalone, specify 1 as the <i>module-number</i>.</p>	<b>set inlinepower management</b> {realtime   class} <i>module-number</i>

**Procedure 14-1 PoE Configuration (continued)**

Step	Task	Command(s)
	Use the <b>clear</b> command to reset the PoE management mode on a specified module back to the default setting of <b>realtime</b> .	<b>clear inlinepower management</b> <i>module-number</i>
7.	<p>(Optional) Configure the allocation mode for system power available for PoE.</p> <ul style="list-style-type: none"> <li>• <b>auto</b> — Available power is distributed evenly to PoE modules based on PoE port count. Any change in available power, due to a change in power supply status or redundancy mode or to the addition or removal of modules, triggers an automatic redistribution of power to the PoE controller on each PoE module.</li> <li>• <b>manual</b> — The power budget for each PoE module is configured manually, using the <b>set inlinepower assigned</b> command.</li> </ul> <p>The configured wattage assignments are used to calculate each module's percentage of total available power. If the total available PoE power changes, a redistribution of available power occurs, applying the calculated percentages.</p> <p>Use this command to reset chassis power allocation to the default mode of <b>auto</b>.</p>	<p><b>set inlinepower mode {auto   manual}</b></p> <p><b>clear inlinepower mode</b></p>
8.	<p>(Only if the <b>set inlinepower mode</b> command is set to <b>manual</b>) Assign specific wattage to a PoE module.</p> <p>If the <b>set inlinepower mode</b> command is set to <b>manual</b>, you must assign power to each PoE module; otherwise, the module ports will not receive power.</p> <p>If the value set with this command is greater than the maximum power percentage specified with the <b>set inlinepower available</b> command, a warning will display in the <b>show inlinepower</b> output. If you execute these parameters, a ratio of assigned power is applied to each module.</p> <p>If your 7100-Series device is a Standalone, specify 1 as the <i>slot-number</i>.</p> <p>Use the <b>clear</b> command to clear the power value manually assigned to one or more modules.</p>	<p><b>set inlinepower assigned</b> <i>power-value</i> <i>slot-number</i></p> <p><b>clear inlinepower assigned</b> [<i>slot-number</i>]</p>

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

## PoE Display Commands

Table 14-3 lists PoE show commands for 7100-Series devices.

**Table 14-3 PoE Show Commands**

Task	Command
Use this command to display PoE properties for a device.	<b>show inlinepower</b>
Use this command to display information about the ports that support PoE: <ul style="list-style-type: none"><li>• Type of PD attached (if specified)</li><li>• Administrative and operational status</li><li>• Priority</li><li>• Class of PD attached</li><li>• Power used by the PD</li></ul>	<b>show port inlinepower</b> [ <i>port-string</i> ]

Refer to the *Extreme Networks 7100-Series CLI Reference* for a description of the output of each command.



## Discovery Protocol Configuration

This chapter provides information about configuring and monitoring discovery protocols on the 7100-Series devices.

For information about...	Refer to page...
<a href="#">How to Use Neighbor Discovery in Your Network</a>	15-1
<a href="#">Understanding Neighbor Discovery</a>	15-2
<a href="#">Configuring LLDP</a>	15-7
<a href="#">Configuring Enterasys Discovery Protocol</a>	15-12
<a href="#">Configuring Cisco Discovery Protocol</a>	15-13

### How to Use Neighbor Discovery in Your Network

Neighbor discovery is the Layer 2 process in which a device identifies and advertises itself to its directly connected neighbors. Extreme Networks devices support the following neighbor discovery protocols:

- Link Layer Discovery Protocol (LLDP) and its extension, LLDP-MED, which is the IEEE 802.1AB standard for neighbor discovery
- Enterasys Discovery Protocol, for discovering Extreme Networks devices
- Cisco Discovery Protocol, for discovering Cisco devices

Neighbor discovery is useful for

- Determining an accurate physical network topology
- Creating an inventory of network devices
- Troubleshooting the network

LLDP, Enterasys Discovery Protocol, and Cisco Discovery Protocol are enabled on Extreme Networks devices by default. Though all three discovery protocols can run simultaneously, LLDP is the preferred protocol.

If a device, attached to a port that has been enabled for neighbor discovery, does not support LLDP but supports Enterasys Discovery Protocol or Cisco Discovery Protocol, then one of those protocols is used instead.



## Understanding Neighbor Discovery

The neighbor discovery protocols support the Layer 2 process of network devices advertising their identities and capabilities on a LAN and discovering that information about their directly connected neighbors. While Enterasys Discovery Protocol and Cisco Discovery Protocol are vendor-specific protocols, LLDP is an industry standard (IEEE 802.1AB), vendor-neutral protocol.

The LLDP-enabled device periodically advertises information about itself (such as management address, capabilities, media-specific configuration information) in an LLDPDU (Link Layer Discovery Protocol Data Unit), which is sent in a single 802.3 Ethernet frame (see [Figure 15-3](#) on page 15-6). An LLDPDU consists of a set of TLV (type, length, and value) attributes. The information, which is extracted and tabulated by an LLDP-enabled device's peers, is recorded in IEEE-defined management information base (MIB) modules, making it possible for the information to be accessed by a network management system using a management protocol such as SNMP. The information is aged to ensure that it is kept up to date. Ports can be configured to send this information, receive this information, or both.

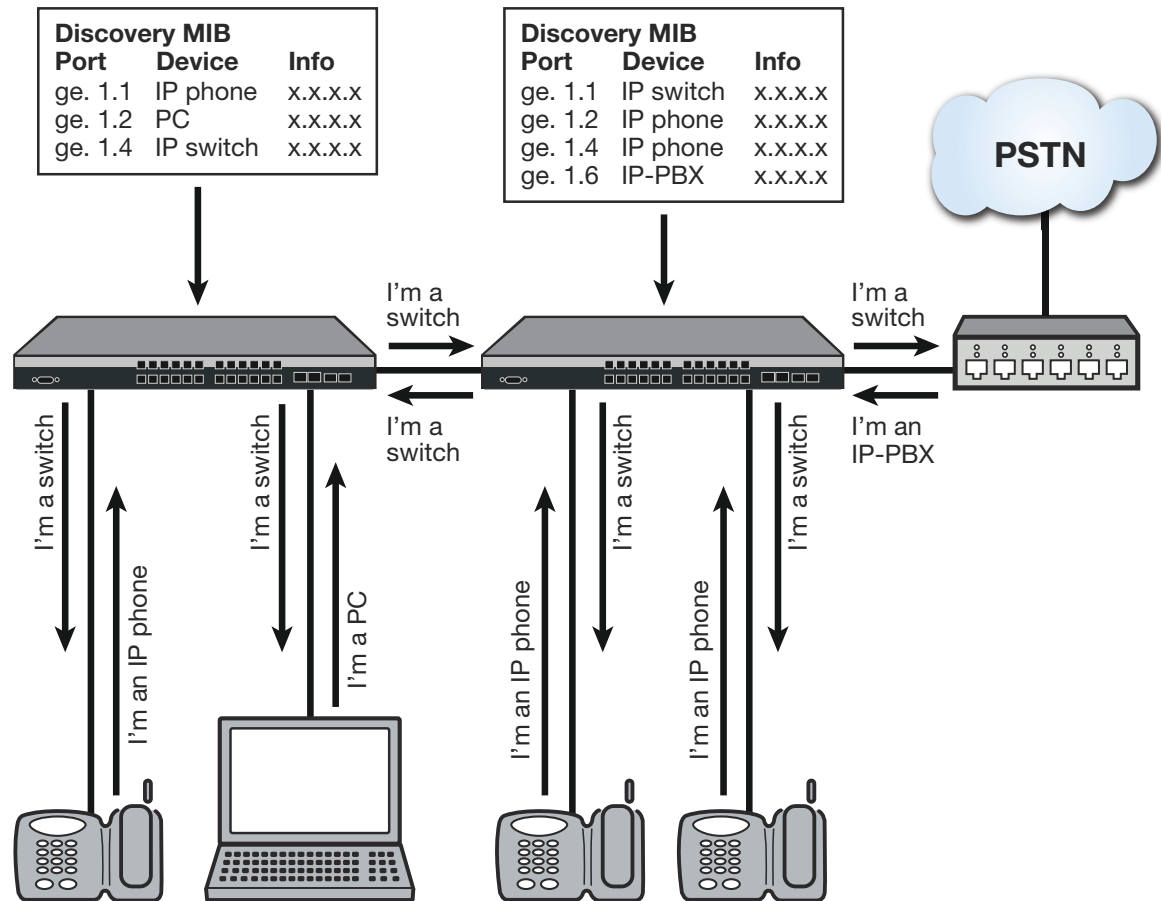
The LLDP agent operates only in an advertising mode, and hence does not support any means for soliciting information or keeping state between two LLDP entities.

LLDP can be used for many advanced features in a VoIP network environment. These features include basic configuration, network policy configuration, location identification (including for Emergency Call Service/E911), Power over Ethernet management, and inventory management.

To fulfill these needs, the standard provides extensions to IEEE 802.1AB that are specific to the requirements of media endpoint devices in an IEEE 802 LAN. Interaction behavior between the media endpoint devices and the LAN infrastructure elements are also described where they are relevant to correct operation or multi-vendor interoperability. Media endpoint devices addressed include, but are not limited to, IP phones, IP voice/media gateways, IP media servers, and IP communication controllers.

The 7100-Series device supports a neighbor warning detection feature which enables protocol checking for a set of potential misconfigurations between this device and the neighbor port.

[Figure 15-1](#) on page 15-3 shows an example of LLDP communication between devices, done via Layer 2 with LLDPDU packets. The communication is only between LLDP-enabled devices — the information is not forwarded to other devices.

**Figure 15-1 Communication between LLDP-enabled Devices**

## LLDP-MED

The LLDP-Media Endpoint Discovery (LLDP-MED) extension of LLDP is defined to share information between media endpoint devices such as IP telephones, media gateways, media servers, and network connectivity devices.

Either LLDP or LLDP-MED, but not both, can be used on an interface between two devices. A switch port uses LLDP-MED when it detects that an LLDP-MED device is connected to it.

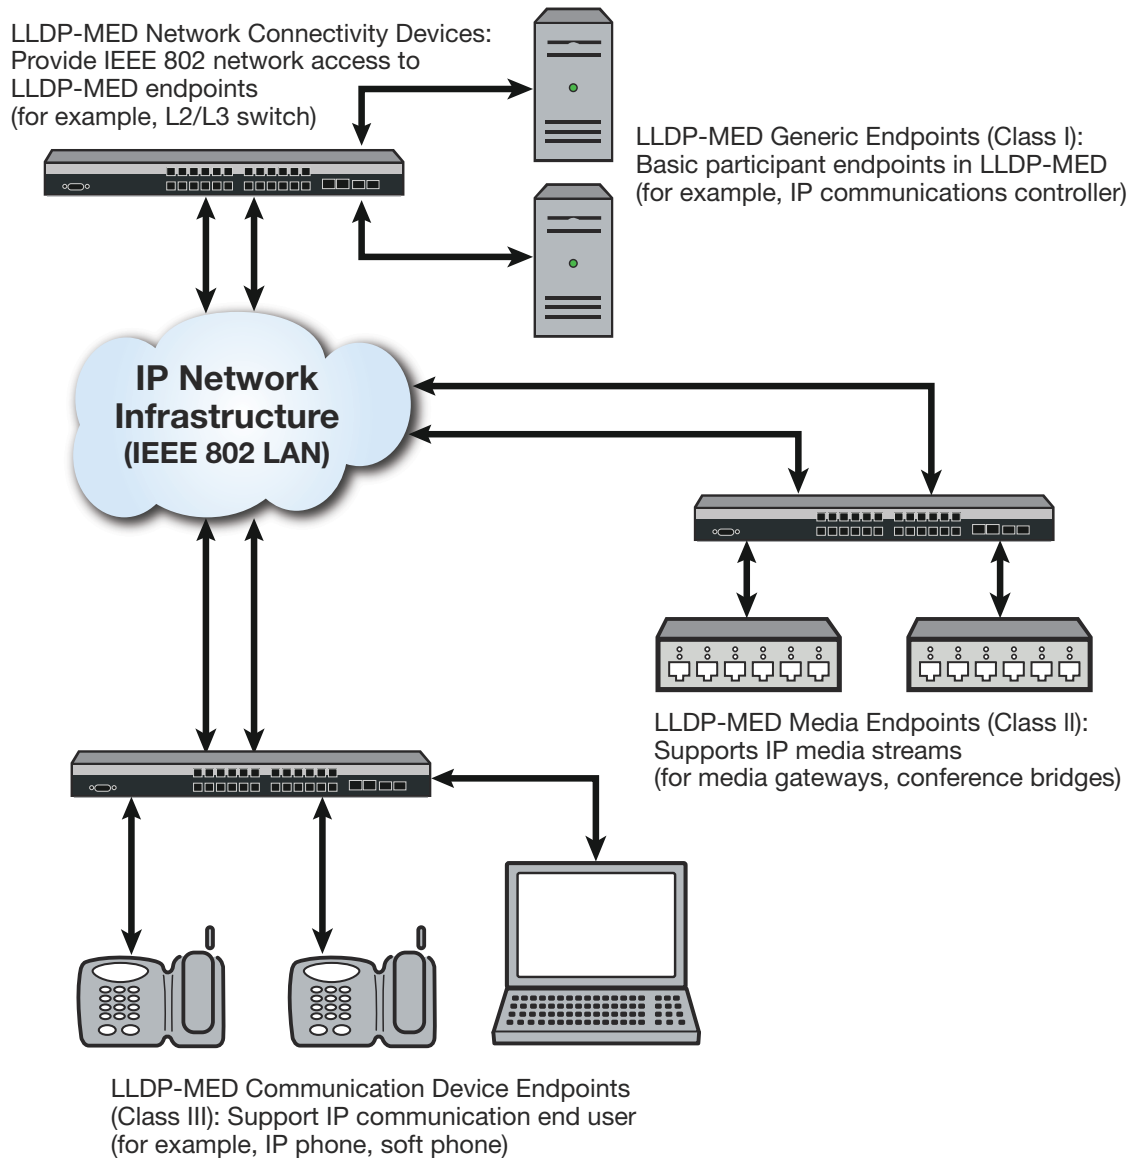
LLDP-MED provides the following benefits:

- Auto discovery of LAN policies, such as VLAN ID, 802.1p priority, and DiffServ codepoint settings, leading to plug-and-play networking.
- Device location and topology discovery, allowing creation of location databases and, in the case of VoIP, provision of E911 services.
- Extended and automated power management of Power over Ethernet endpoints
- Inventory management, allowing network administrators to track their network devices and to determine their characteristics, such as manufacturer, software and hardware versions, and serial or asset numbers.

There are two primary LLDP-MED device types (as shown in [Figure 15-2](#) on page 15-5):

- Network connectivity devices, which are LAN access devices such as LAN switch/router, bridge, repeater, wireless access point, or any device that supports the IEEE 802.1AB and MED extensions defined by the standard and can relay IEEE 802 frames via any method.

- Endpoint devices, which have three defined sub-types or classes:
  - LLDP-MED Generic Endpoint (Class I) — All endpoint products that, while requiring the base LLDP discovery services defined in the standard, do not support IP media or act as an end-user communication device, such as IP communications controllers, other communication-related servers, or any device requiring basic services. Discovery services defined in this class include LAN configuration, device location, network policy, power management, and inventory management.
  - LLDP-MED Media Endpoint (Class II) — All endpoint products that have IP media capabilities but that may not be associated with a particular end user, such as voice/media gateways, conference bridges, and media servers. Capabilities include all of the capabilities defined for Generic Endpoint (Class I) and are extended to include aspects related to media streaming. Discovery services defined in this class include media type specific network layer policy discovery.
  - LLDP-MED Communication Endpoint (Class III) — All endpoint products that act as an endpoint user communication device supporting IP media. Capabilities include all of the capabilities defined for the Generic Endpoint (Class I) and Media Endpoint (Class II) devices and are extended to include aspects related to end user devices, such as IP phones, PC-based soft phones, and other communication devices that directly support the end user.



**Figure 15-2 LLDP-MED**

## LLDPDU Frames

As shown in [Figure 15-3](#), each LLDPDU frame contains the following mandatory TLVs:

- Chassis ID — The chassis identification for the device that transmitted the LLDP packet.
- Port ID — The identification of the specific port that transmitted the LLDP packet. The receiving LLDP agent joins the chassis ID and the port ID to correspond to the entity connected to the port where the packet was received.
- Time to Live — The length of time that information contained in the receive LLDP packet will be valid.
- End of LLDPDU — Indicates the final TLV of the LLDPDU frame.

**Figure 15-3 Frame Format**



M = Mandatory TLV (required for all LLDPDUs)

Each LLDPDU frame can also contain the following optional TLVs:

- Port Description — The port from which the LLDP agent transmitted the frame.
- System Name — The system’s administratively assigned name.
- System Description — Includes the system’s name, hardware version, OS level, and networking software version.
- System Capabilities — A bitmap that defines the primary functions of the system. The currently defined capabilities include, among other things, WLAN access point, router, and telephone.
- Management Address — The IP or MAC address associated with the local LLDP agent that may be used to reach higher layer entities.

An LLDPDU frame can also contain the following extension TLVs:

- 802.1 VLAN extension TLVs describe attributes associated with VLANs:
  - Port VLAN ID — Allows a bridge port to advertise the port’s VLAN identifier (PVID) that will be associated with untagged or priority tagged frames it receives.
  - Port & Protocol VLAN ID — Allows a bridge to advertise whether it supports protocol VLANs and, if so, what VLAN IDs these protocols will be associated with.
  - VLAN Name — Allows a bridge to advertise the textual name of any VLAN with which it is configured.
  - Protocol Identity — Allows a bridge to advertise the particular protocols that are accessible through its port.
- 802.3 LAN interface extensions TLVs describe attributes associated with the operation of an 802.3 LAN interface:
  - MAC/PHY Configuration/Status — Advertises the bit-rate and duplex capability of the sending 802.3 node, the current duplex and bit-rating of the sending 802.3 node, and whether these settings were the result of auto-negotiation during link initiation or manual override.
  - Power-Via-MDI — Advertises the power-via-MDI capabilities of the sending 802.3 node.
  - Link-Aggregation — Advertises whether the link is capable of being aggregated, whether it is currently in an aggregation, and, if it is in an aggregation, the port of the aggregation.

- Maximum Frame Size — Advertises the maximum supported 802.3 frame size of the sending station.
- LLDP-MED extension TLVs:
  - Capabilities — Indicates the network connectivity device’s capabilities.
  - Network Policy — Used to configure tagged/untagged VLAN ID/L2 priority/DSCP on LLDP-MED endpoints (for example, IP phones).
  - Location Identification — Provides the location identifier information to communication endpoint devices, based on the configuration of the network connectivity device it is connected to.
  - Extended Power via MDI — Enables advanced power management between LLDP-MED endpoints and network connectivity devices.
  - Inventory Management — Includes hardware revision, firmware revision, software revision, serial number, manufacturer name, model name, and asset ID.

Some TLVs support multiple subtypes. For example, Port ID is sent as an ifName (e.g., ge.1.1) between Extreme Networks devices, but when an LLDP-MED endpoint is detected on a port, that TLV subtype changes to a network address (MAC address), and other MED TLVs are sent, as defined by the MED spec.

## Neighbor Warning Detection

The 7100-Series device supports a neighbor warning detection feature which enables checking for a set of potential misconfigurations between this device and the neighbor port. Warning types are configurable on a per port basis. Generated warnings can be displayed using a `warnings show` command. Warning types support the detection of neighbor differences in:

- Speed and duplex
- Power class
- MTU
- LACP status
- PFC status

## Configuring LLDP

### LLDP Configuration Commands

[Table 15-1](#) lists LLDP configuration commands. The table indicates which commands are device specific.

**Table 15-1 LLDP Configuration Commands**

Task	Command
Set the time, in seconds, between successive LLDP frame transmissions initiated by changes in the LLDP local system information. Default value is 30 seconds.	<code>set lldp tx-interval <i>frequency</i></code>
Set the number of LLDP PDU packets sent when entering fast transmission state.	<code>set lldp tx-fast-count <i>count</i></code>

**Table 15-1 LLDP Configuration Commands (continued)**

Task	Command
Set the frequency of LLDP PDU transmissions while in fast transmission state.	<b>set lldp tx-fast-interval</b> <i>frequency</i>
Set the time-to-live value used in LLDP frames sent by this device. The time-to-live for LLDPDU data is calculated by multiplying the transmit interval by the hold multiplier. The default value is 4.	<b>set lldp hold-multiplier</b> <i>multiplier-val</i>
Set the minimum interval between LLDP notifications sent by this device. LLDP notifications are sent when a remote system change has been detected. The default value is 5 seconds.	<b>set lldp trap-interval</b> <i>frequency</i>
Set the number of fast start LLDPDUs to be sent when an LLDP-MED endpoint device, such as a phone, is detected. Network connectivity devices transmit only LLDP TLVs in LLDPDUs until they detect that an LLDP-MED endpoint device has connected to a port. At that point, the network connectivity device starts sending LLDP-MED TLVs at a fast start rate on that port. The default value is 3.	<b>set lldp med-fast-repeat</b> <i>count</i>
Enable or disable transmitting and processing received LLDPDUs on a port or range of ports.	<b>set lldp port status</b> { <b>tx-enable</b>   <b>rx-enable</b>   <b>both</b>   <b>disable</b> } <i>port-string</i>
Enable or disable sending LLDP traps when a remote system change is detected.	<b>set lldp port trap</b> { <b>enable</b>   <b>disable</b> } <i>port-string</i>
Enable or disable sending an LLDP-MED trap when a change in the topology has been sensed on the port (that is, a remote endpoint device has been attached or removed from the port).	<b>set lldp port med-trap</b> { <b>enable</b>   <b>disable</b> } <i>port-string</i>
Configure LLDP-MED location information on a port or range of ports. Currently, only Emergency Call Services (ECS) Emergency Location Identification Number (ELIN) is supported. ELIN is a special phone number used to indicate location, and is assigned and associated with small geographies in the organization. It is one of the forms of identification that the location identification TLV provides.	<b>set lldp port location-info elin</b> <i>elin-string</i> <i>port-string</i>
Select the optional LLDP and LLDP-MED TLVs to be transmitted in LLDPDUs by the specified port or ports.	<b>set lldp port tx-tlv</b> {[ <b>all</b> ]   [ <b>port-desc</b> ] [ <b>sys-name</b> ] [ <b>sys-desc</b> ] [ <b>sys-cap</b> ] [ <b>mgmt-addr</b> ] [ <b>vlan-id</b> ] [ <b>stp</b> ] [ <b>lACP</b> ] [ <b>gvrp</b> ] [ <b>mac-phy</b> ] [ <b>poe</b> ] [ <b>link-aggr</b> ] [ <b>max-frame</b> ] [ <b>med-cap</b> ] [ <b>med-pol</b> ] [ <b>med-loc</b> ] [ <b>med-poe</b> ] [ <b>enhanced-trans-config</b> ] [ <b>enhanced-trans-rec</b> ] [ <b>priority-flowctrl</b> ]} <i>port-string</i>
Configure network policy for a set of applications on a port or range of ports. The policies configured with this command are sent in LLDPDUs as LLDP-MED Network Policy TLVs. Multiple Network Policy TLVs can be sent in a single LLDPDU.	<b>set lldp port network-policy</b> { <b>all</b>   <b>voice</b>   <b>voice-signaling</b>   <b>guest-voice</b>   <b>guest-voice-signaling</b>   <b>softphone-voice</b>   <b>video-conferencing</b>   <b>streaming-video</b>   <b>video-signaling</b> } [ <b>state</b> { <b>enable</b>   <b>disable</b> }] [ <b>tag</b> { <b>tagged</b>   <b>untagged</b> }] [ <b>vid</b> { <i>vlan-id</i>   <b>dot1p</b> }] [ <b>cos</b> <i>cos-value</i> ] [ <b>dscp</b> <i>dscp-value</i> ] <i>port-string</i>
Return LLDP parameters to their default values.	<b>clear lldp</b> { <b>all</b>   <b>tx-interval</b>   <b>hold-multiplier</b>   <b>trap-interval</b>   <b>med-fast-repeat</b> }

**Table 15-1 LLDP Configuration Commands (continued)**

Task	Command
Return the port status to the default value of both (both transmitting and processing received LLDPDUs are enabled).	<b>clear lldp port status</b> <i>port-string</i>
Return the port LLDP trap setting to the default value of disabled.	<b>clear lldp port trap</b> <i>port-string</i>
Return the port LLDP-MED trap setting to the default value of disabled.	<b>clear lldp port med-trap</b> <i>port-string</i>
Return the port ECS ELIN location setting to the default value of null.	<b>clear lldp port location-info elin</b> <i>port-string</i>
Return network policy for a set of applications on a port or range of ports to default values.	<b>clear lldp port network-policy</b> {all   voice   voice-signaling   guest-voice   guest-voice-signaling   softphone-voice   video-conferencing   streaming-video   video-signaling} {[state] [tag] [vid] [cos] [dscp]} <i>port-string</i>
Clear the optional LLDP and LLDP-MED TLVs to be transmitted in LLDPDUs by the specified port or ports to the default value of disabled.	<b>clear lldp port tx-tlv</b> {[all]   [port-desc] [sys-name] [sys-desc] [sys-cap] [mgmtaddr] [vlan-id] [stp] [lacp] [gvrp] [mac-phy] [poe] [link-aggr] [max-frame] [medcap] [med-pol] [med-loc] [med-poe]} <i>port-string</i>

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

## Basic LLDP Configuration

[Procedure 15-1](#) describes the basic steps to configure LLDP on Extreme Networks 7100-Series devices.

### Procedure 15-1 Configuring LLDP (Extreme Networks 7100-Series)

Step	Task	Command(s)
1.	Configure global system LLDP parameters.	<b>set lldp tx-interval</b> <b>set lldp hold-multiplier</b> <b>set lldp trap-interval</b> <b>set lldp med-fast-repeat</b> <b>clear lldp</b>
2.	Enable/disable specific ports to: <ul style="list-style-type: none"> <li>Transmit and process received LLDPDUs</li> <li>Send LLDP traps</li> <li>Send LLDP-MED traps</li> </ul>	<ul style="list-style-type: none"> <li><b>set/clear lldp port status</b></li> <li><b>set/clear lldp port trap</b></li> <li><b>set/clear lldp port med-trap</b></li> </ul>
3.	Configure an ECS ELIN value for specific ports.	<b>set/clear lldp port location-info</b>



**Procedure 15-1 Configuring LLDP (Extreme Networks 7100-Series) (continued)**

Step	Task	Command(s)
4.	Configure Network Policy TLVs for specific ports.	<b>set/clear lldp port network-policy</b>
5.	Configure which optional TLVs should be sent by specific ports. For example, if you configured an ECS ELIN and/or Network Policy TLVs, you must enable those optional TLVs to be transmitted on the specific ports.	<b>set/clear lldp tx-tlv</b>

**Example LLDP Configuration: Time to Live**

This example sets the transmit interval to 20 seconds and the hold multiplier to 5, which will configure a time-to-live of 100 to be used in the TTL field in the LLDPDU header.

```
7100 Chassis(rw)->set lldp tx-interval 20
7100 Chassis(rw)->set lldp hold-multiplier 5
```

**Example LLDP Configuration: Location Information**

On an 7100-Series device, after you configure a location information value, you must also configure the port to send the Location Information TLV with the **set lldp port tx-tlv** command. This example configures the ELIN identifier 5551234567 on ports ge.1.1 through ge.1.6 and then configures the ports to send the Location Information TLV.

```
7100 Chassis(rw)->set lldp port location-info 5551234567 ge.1.1-6
7100 Chassis(rw)->set lldp port tx-tlv med-loc ge.1.1-6
```

**LLDP Display Commands**

[Table 15-2](#) lists LLDP show commands. The table indicates which commands are device specific.

**Table 15-2 LLDP Show Commands**

Task	Command
Display LLDP configuration information.	<b>show lldp</b>
Display the LLDP status of one or more ports.	<b>show lldp port status</b> [ <i>port-string</i> ]
Display the ports that are enabled to send an LLDP notification when a remote system change has been detected or an LLDP-MED notification when a change in the topology has been sensed.	<b>show lldp port trap</b> [ <i>port-string</i> ]
Display information about which optional TLVs have been configured to be transmitted on ports.	<b>show lldp port tx-tlv</b> [ <i>data-center-bridging</i> ] [ <i>port-string</i> ]
Display configured location information for one or more ports.	<b>show lldp port location-info</b> [ <i>port-string</i> ]
Display the local system information stored for one or more ports.	<b>show lldp port local-info</b> [ <i>port-string</i> ]
Display the remote system information stored for a remote device connected to a local port.	<b>show lldp port remote-info</b> [ <i>port-string</i> ]

**Table 15-2 LLDP Show Commands (continued)**

Task	Command
Display LLDP port network policy configuration information.	<b>show lldp port network policy</b> {all   voice   voice-signaling   guest-voice   guestvoice-signaling   software-voice   video-conferencing   streaming-video   videosignaling} [ <i>port-string</i> ]

Refer to the *Extreme Networks 7100-Series CLI Reference* for a description of the output of each command.

## Configuring Neighbor Warning Detection

Table 15-3 lists neighbor warning detection configuration and display commands.

**Table 15-3 Neighbor Warning Detection Commands**

Task	Command
Enable protocol checking for port mis-configuration (warning detection) with its neighbor	<b>set neighbors warning-detection</b> <i>warning-type</i> [ <i>port-string</i> ] {enable   disable}
Display neighbors with warnings due to enabled warning detection hits on the port.	<b>show neighbors warnings</b> [ <i>warning-type</i> ] [ <i>port-string</i> ]
Display the status of each neighbors warning type for all or specified ports.	<b>show neighbors warning-detection</b> [ <i>port-string</i> ]

## Configuring Enterasys Discovery Protocol

### Enterasys Discovery Protocol Configuration Commands

Table 15-4 lists Enterasys Discovery Protocol configuration commands.

**Table 15-4 Enterasys Discovery Protocol Configuration Commands**

Task	Command
Enable or disable the Enterasys Discovery Protocol on one or more ports.	<b>set cdp state</b> {auto   disable   enable} [ <i>port-string</i> ]
Set a global Enterasys Discovery Protocol authentication code.	<b>set cdp auth</b> <i>auth-code</i>
Set the message interval frequency (in seconds) of the Enterasys Discovery Protocol.	<b>set cdp interval</b> <i>frequency</i>
Set the hold time value for Enterasys Discovery Protocol configuration messages.	<b>set cdp hold-time</b> <i>hold-time</i>
Reset Enterasys Discovery Protocol settings to defaults.	<b>clear cdp</b> {[state] [port-state <i>port-string</i> ] [interval] [hold-time] [auth-code]}

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

### Example Enterasys Discovery Protocol Configuration

This example shows how to globally enable CDP:

```
7100 Chassis(rw)->set cdp state enable
```

This example shows how to enable the CDP for port ge.1.2:

```
7100 Chassis(rw)->set cdp state enable ge.1.2
```

This example shows how to disable the CDP for port ge.1.2:

```
7100 Chassis(rw)->set cdp state disable ge.1.2
```

## Enterasys Discovery Protocol Show Commands

Table 15-5 lists Enterasys Discovery Protocol show commands.

**Table 15-5 Enterasys Discovery Protocol Show Commands**

Task	Command
Display the status of the CDP discovery protocol and message interval on one or more ports.	<b>show cdp</b> [ <i>port-string</i> ]
Display Network Neighbor Discovery information from all supported discovery protocols.	<b>show neighbors</b> [ <i>port-string</i> ]

Refer to the *Extreme Networks 7100-Series CLI Reference* for a description of the output of each command.

## Configuring Cisco Discovery Protocol

### Cisco Discovery Protocol Configuration Commands

Table 15-6 lists Cisco Discovery Protocol configuration commands.

**Table 15-6 Cisco Discovery Protocol Configuration Commands**

Task	Command
Enable or disable Cisco Discovery Protocol globally on the device.	<b>set ciscodp status</b> { <i>auto</i>   <i>enable</i>   <i>disable</i> }
Set the number of seconds between Cisco Discovery Protocol PDU transmissions.	<b>set ciscodp timer</b> <i>time</i>
Set the time to live (TTL) for Cisco Discovery Protocol PDUs. This is the amount of time (in seconds) neighboring devices will hold PDU transmissions from the sending device.	<b>set ciscodp holdtime</b> <i>time</i>
Set the status, voice VLAN, extended trust mode, and CoS priority for untrusted traffic for the Cisco Discovery Protocol on one or more ports.	<b>set ciscodp port</b> { [ <i>status</i> { <i>disable</i>   <i>enable</i> }] [ <i>vvid</i> {< <i>vlan-id</i> >   <i>none</i>   <i>dot1p</i>   <i>untagged</i> }] [ <i>trust-ext</i> { <i>trusted</i>   <i>untrusted</i> }] [ <i>cos-ext value</i> ] } < <i>port-string</i> >
Clear the Cisco Discovery Protocol back to the default values.	<b>clear ciscodp</b> { [ <i>status</i>   <i>timer</i>   <i>holdtime</i>   <i>port</i> { <i>status</i>   <i>vvid</i>   <i>trust-ext</i>   <i>cos-ext</i> }] } < <i>port-string</i> >

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

### Example Cisco Discovery Protocol Configuration

This example shows how to enable Cisco Discovery Protocol on the device:

```
7100 Chassis(rw)->set ciscodp status enable
```

### Cisco Discovery Protocol Show Commands

Table 15-7 lists Cisco Discovery Protocol show commands.

**Table 15-7 Cisco Discovery Protocol Show Commands**

<b>Task</b>	<b>Command</b>
Display global Cisco Discovery Protocol information.	<b>show cisco</b> <b>dp</b>
Display summary information about the Cisco Discovery Protocol on one or more ports.	<b>show cisco</b> <b>dp port info</b> [ <i>port-string</i> ]
Display Network Neighbor Discovery information from all supported discovery protocols.	<b>show neighbors</b> [ <i>port-string</i> ]

Refer to the *Extreme Networks 7100-Series CLI Reference* for a description of the output of each command.

## Data Center Bridging Configuration

This chapter provides information about configuring and monitoring Data Center Bridging (DCB) protocols on 7100-Series devices.

For information about...	Refer to page...
<a href="#">How to Use Data Center Bridging in Your Network</a>	16-1
<a href="#">Implementing Data Center Bridging</a>	16-2
<a href="#">Enhanced Transmission Selection Configuration</a>	16-2
<a href="#">Priority-Based Flow Control Configuration</a>	16-4
<a href="#">Application Priority Configuration</a>	16-5
<a href="#">Congestion Notification (CN) Configuration</a>	16-5
<a href="#">Configuring Data Center Bridging</a>	16-14
<a href="#">Terms and Definitions</a>	16-17

### How to Use Data Center Bridging in Your Network

Data Center Bridging (DCB) enhances Ethernet technology by enabling the convergence of various applications in data centers (such as Local Area Networks (LAN), Storage Area Networks (SAN), and advanced application High Performance Computing (HPC)) onto a single interconnect technology, by providing enhancements to existing 802.1 bridge specifications. Existing high-performance data centers typically comprise multiple application-specific networks that run on different link layer technologies, such as Fibre Channel for storage, InfiniBand for high-performance computing, and Ethernet for network management and LAN connectivity. Data Center Bridging enables 802.1 bridges to be used for the deployment of a converged network where all applications can be run over a single physical infrastructure.

The current Data Center Bridging implementation consists of:

- **Enhanced Transmission Selection (ETS):** Provides a common management framework for assignment of bandwidth to 802.1p CoS-based traffic classes (IEEE 802.1Qaz).
- **Congestion Notification (CN):** Allows a device to detect congestion on an egress transmit queue and send a message back to the source to back off the traffic rate to alleviate the congestion (IEEE 802.1Q-2011).
- **Priority-based Flow Control (PFC):** Provides a link level flow control mechanism that can be controlled independently for each Class of Service (CoS), as defined by 802.1p. The goal of this mechanism is to ensure zero loss under congestion in Data Center Bridging networks (IEEE 802.1Qbb).

- **Application Priority:** Provides for the advertisement to the link peer of a preferred priority to be applied to frames carrying application-specific traffic.

The base control protocol utilized in Data Center Bridging is the Data Center Bridging Exchange (DCBX) protocol. DCBX can be used by a device: to detect peer device capabilities, to detect mis-configuration of a feature between the peers on a link, to perform configuration of DCB features on the link peer.

Enhanced Transmission Selection, priority-based flow control, congestion notification, and Application Priority protocols utilize DCBX. DCBX uses LLDP to exchange attributes between two linked peers. LLDP is unidirectional and advertises connectivity and management information about the local station to adjacent stations on the same IEEE 802 LAN. DCBX state machines are invoked when the remote MIB changes and a DCBX TLV is present.

## Implementing Data Center Bridging

This chapter describes how to configure DCB features on Extreme Networks platforms.

- Configure ports for Enhanced Transmission Selection specifying the ETS groups and the bandwidth allocation assigned to each ETS group and optionally advertising the ETS configuration using the LLDP-DCB Enhanced Transmission Configuration and Enhanced Transmission Recommendation TLVs.
- Configure ports for priority-based flow control by enabling priority-based flow control on a port for specified 802.1p priorities and advertising the PFC configuration to peers using the LLDP Priority Flow Control TLV.
- Configure ports to advertise an Application Priority setting to the LLDP link peer for the specified application and advertising the Application Priority configuration using the LLDP Application Priority TLV.
- Configure switch ports for congestion notification.

You configure DCB features separately and independently. Together the four features implement DCB.

## Enhanced Transmission Selection Configuration

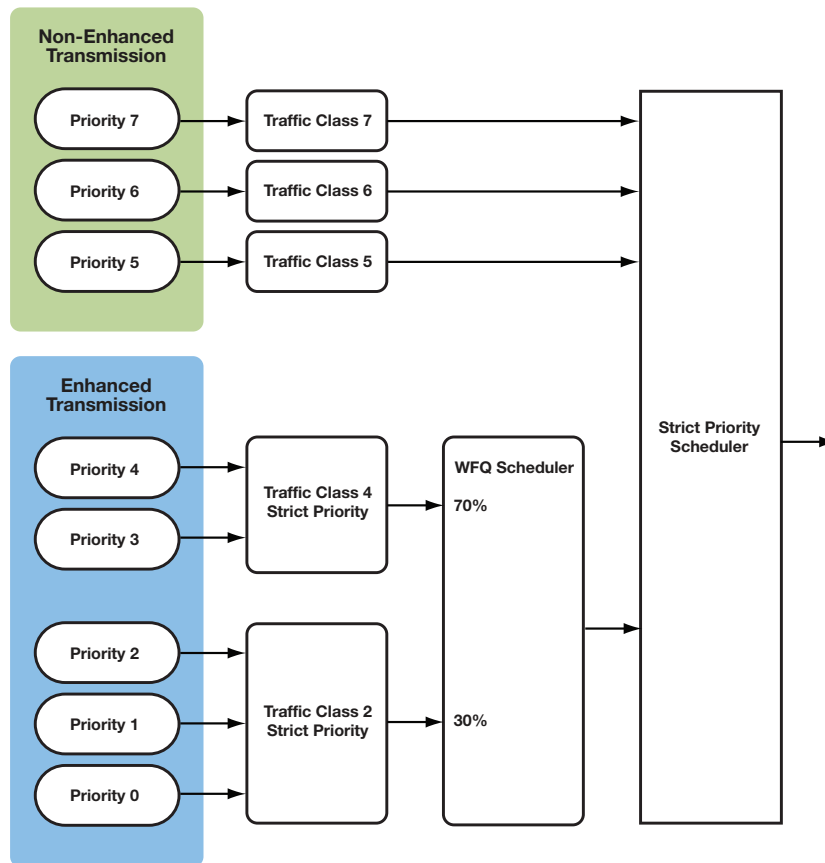
Enhanced Transmission Selection (ETS) queuing provides for configuring two or more traffic class queues (transmit queue (TxQ)) to be allocated for bandwidth that will not be serviced until all non-ETS queues are empty. The firmware services non-ETS selection queues first using strict priority, based upon the priority assigned to the queue. Enhanced Transmission Selection queue contents are forwarded to a fair queue scheduler on a strict priority basis. The fair queue scheduler distributes the remaining bandwidth, after all non-Enhanced Transmission Selection queues are empty, based upon the bandwidth allocation configured for the Enhanced Transmission Selection queues.



**Note:** Enhanced Transmission Selection queuing is restricted to configurable queues. 7100-Series modules support both configurable and non-configurable queues. Non-configurable queues are Low Latency Queues (LLQ). LLQs are labeled LLQ in the **show cos port-config** command display. See “[Low Latency Queuing](#)” on page 40-4 for LLQ details.

Enhanced Transmission Selection traffic classes (TxQs) and bandwidth allocation are configured using the **set cos port-config txq** command.

[Figure 16-1](#) presents an Enhanced Transmission Selection queuing example. Priority queues 7, 6, and 5 are assigned traffic classes 7, 6, and 5, respectively. These non-ETS queues are serviced first by the strict priority scheduler based upon priority.

**Figure 16-1 Enhanced Transmission Selection (ETS) Queuing**

802.1 priorities 0 - 4 are configured for Enhanced Transmission Selection queuing. Priorities 4 and 3 are assigned to traffic class 4. Priorities 0, 1, and 2 are assigned to traffic class 2. If all non-ETS queues are empty and there is remaining bandwidth, traffic classes 4 and 2 will be serviced using weighted fair queue scheduling. Based upon Enhanced Transmission Selection bandwidth allocation, the weighted fair queue scheduler will service traffic class 4 at 70 percent and traffic class 2 at 30 percent of remaining bandwidth. Within each traffic class group (4 and 2 in this example), each priority is serviced based on a strict priority scheduler.



**Note:** The 7100-Series supports up to two ETS groups (traffic classes).

This example shows how to create a CoS transmit queue port group entry named **testTxq** with a port group ID of 2 and a port type ID of 0, assign ETS groups to an 8 queue device, followed by allocation of ETS bandwidth to the assigned groups. Using the **enhanced-groups** option of the **set cos port-config txq** command, ETS group to queue assignment is:

- Group 2 to queues 0, 1, and 2
- Group 4 to queues 3 and 4

Using the **enhanced-percentage** option of the **set cos port-config txq** command, the assigned ETS bandwidth allocation is:

- 30 percent to group 2
- 70 percent to group 4



```
7100 Chassis(rw)->set cos port-config txq 2.0 name testTxq enhanced-groups
2,2,2,4,4,0,0,0 enhanced-percentage 0,30,0,70,0,0,0,0
```

Use the Enhanced Transmission Selection configuration TLV to advertise this ETS configuration to the peer:

```
7100 Chassis(rw)->set lldp port tx-tlv enhanced-trans-config
```

Use the Enhanced Transmission Selection recommendation TLV to recommend that the peer use this ETS configuration:

```
7100 Chassis(rw)->set lldp port tx-tlv enhanced-trans-rec
```

## Priority-Based Flow Control Configuration

Priority-based Flow Control (PFC), as defined by 802.1Qbb, functions similarly to 802.3 PAUSE, but allows for per-priority pausing instead of per-port pausing. This is achieved by pausing the egress queues to which the priorities are mapped.

Traffic congestion is determined on a per-queue basis. When the ingress queue that maps to a particular 802.1p priority reaches a certain (non-configurable) threshold, a priority-based flow control message is sent to the peer to pause transmission of traffic tagged with that priority. Once queue buffer levels return to normal, the ingress queue stops sending priority-based flow control messages, and the peer no longer pauses traffic tagged with that priority.

Priority-based flow control is defined only for a pair of full duplexed MAC devices connected by one point-to-point link. An egress queue is paused by a switch when it receives a message from its peer on the other end of the link that priority-tagged frames mapped to the queue should be paused.

Flows are paused by egress queue, not by priority. If non-PFC priorities are mapped to the same egress queue as PFC priorities, the non-PFC priority data will be paused along with the PFC priority data. For example, if priorities 3 and 6 are set to egress on queue 5 and only priority 3 is enabled for priority-based flow control, when priority 3 is paused, queue 5 will stop transmission and priority 6 will be paused as well.

A buffer allowance is provided for storing incoming traffic between the time a priority-based flow control frame is sent back to the peer and the peer pauses the PFC priority traffic, preventing data loss when initiating priority-based flow control. This buffer amount is called link delay allowance, and is measured in number of bits.

The default priority used primarily for control traffic should not be set for priority-based flow control. Untagged packets bypass egress queues and transmit out with the default priority traffic.



**Note:** Priority-based flow control and 802.3 PAUSE are mutually exclusive. Enabling one feature automatically disables the other. Enabling PFC both clears PAUSE advertisement and prohibits PAUSE advertisement while PFC is enabled.

Use the **set dcb pfc** command to enable priority-based flow control on a port, specifying the port and 802.1p priority.

You can configure LLDP to add the LLDP-DCB Priority-Flowctrl TLV to LLDP PDUs transmitted on a port using the **set lldp port tx-tlv priority-flowctrl** command.

This example shows how to enable priority-based flow control for priority 5 on port **ge.1.2** and advertise the setting to its peer:

```
7100 Chassis(rw)->set dcb pfc ge.1.2 5 enable
7100 Chassis(rw)->set lldp port tx-tlv priority-flowctrl ge.1.2
```

## Application Priority Configuration

Application Priority advertises to the peer a preferred priority for frames carrying application-specific traffic. Applications are defined by protocol (Ethernet, TCP, UDP, or Layer 4 port) and protocol ID. Priority tagging is performed by the peer, not by the device advertising the Application Priority. The peer receiving the Application Priority TLV tags its traffic to the advertised priority. Application Priority works with Enhanced Transmission Selection and priority-based flow control in that tagged protocol-specific traffic for the specified priority enforces Enhanced Transmission Selection and priority-based flow control behaviors on the traffic.

For example, Application Priority could advertise to its peer that all iSCSI traffic be tagged with priority 5. In this case, the TCP protocol on well known port 3260, as assigned to iSCSI by IANA, is specified. The peer receiving the Application Priority TLV will tag iSCSI traffic for priority 5 in the Priority Code Point (PCP) tag of its 802.1Q header. Both Enhanced Transmission Selection and priority-based flow control will use this priority 5 setting if each functionality is configured to do so.



**Note:** The Application Priority feature requires that the peer supports the LLDP willing bit and the willing bit is enabled for priority tagging to occur on the peer. Extreme Networks switches do not currently support the LLDP willing bit. Extreme Networks switches can advertise to the peer a preferred priority for frames carrying application-specific traffic, but an Extreme Networks switch peer will not perform Application Priority tagging.

Use the **set lldp port tx-tlv application-pri** command to enable the advertising of the Application Priority configuration to the link peer using the LLDP-DCB Application Priority TLV.

This example shows how to advertise to the peer priority 4, for the UDP service type dpkeyserver with well known port 1780, for port ge.1.2 and to enable the sending of LLDP-DCB Application Priority TLVs from that port:

```
7100 Chassis(rw)->set dcb appPri ge.1.2 protocol udp protocol-id 1780 priority 4
7100 Chassis(rw)->set lldp port tx-tlv application-pri ge.1.2
```

## Congestion Notification (CN) Configuration

Congestion Notification (CN), as defined in IEEE 802.1Q-2011 allows a device to detect congestion at a switch congestion point (egress transmit queue) and transmit a Congestion Notification Message (CNM) PDU back to the reaction point (flow source). The reaction point backs off the traffic rate to alleviate the congestion. Congestion notification supports long lived data flows in a network with delay due to limited bandwidth. It allows for applications that are latency-or-loss-sensitive to run over Ethernet technologies experiencing egress transmit queue congestion.

As the use of Ethernet technologies in the data center expands, prevention of packet loss by some applications becomes more critical. Congestion notification was created to:

- Allow the monitoring of Congestion Controlled Flows (CCFs)
- Detect congestion
- Notify the source to lower the transmit rate for the offending congestion controlled flow.

For congestion notification to work, it must be supported at all egress queues for each switch that is in the path from the source to the destination. Congestion notification is applied to an egress queue by configuring an 802.1p value as a Congestion Notification Priority Value (CNPV) that is mapped to the transmit queue using CoS. This collection of egress queues configured for congestion notification make up the Congestion Notification Domain (CND).

Each transmit queue that has been configured for congestion notification is monitored to detect congestion. When congestion is detected, a CNM PDU is generated at the congestion point and sent back to the source with the details of the queue and flow that triggered the message. The source can then use this information to back off the transmission rate for the application that triggered the CNM PDU.

SNMP supports the congestion notification IEEE Standard MIB: IEEE8021-CN-MIB

Figure 16-2 provides a Congestion Notification overview.

**Figure 16-2 Congestion Notification Overview**

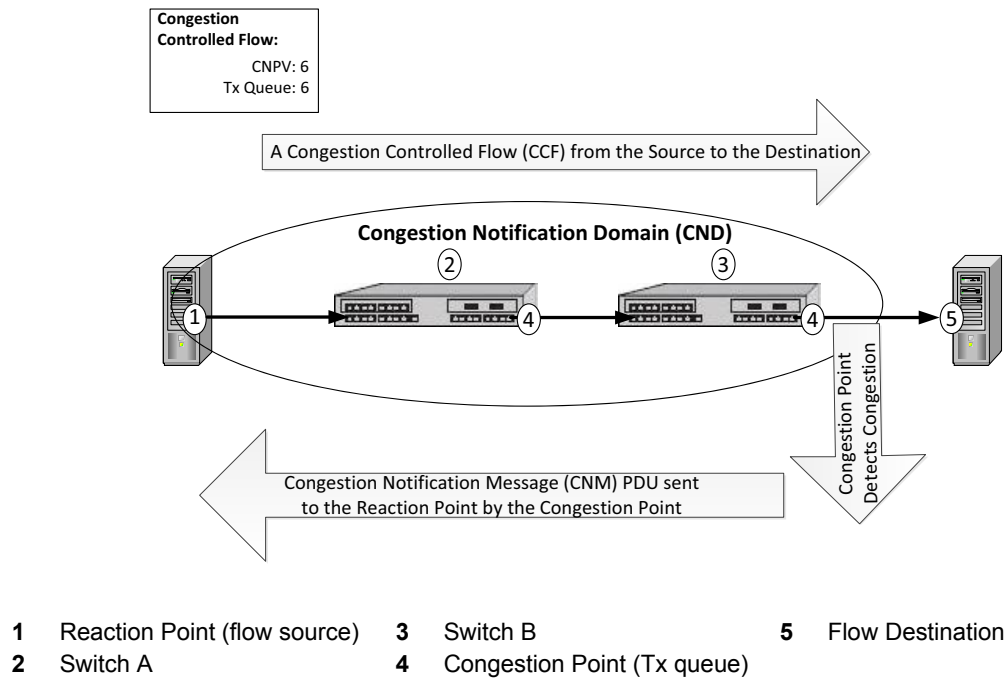


Figure 16-2 identifies the congestion notification reaction point (callout 1). A reaction point is the source of a congestion controlled flow. A congestion controlled flow consists of frames, all with the same CNPV, and all assigned to a single transmit flow queue in the originating end station. A CNPV is an 802.1p priority mapped to a congestion notification egress queue of each device in the flow.

The reaction point is:

- Capable of optionally adding a CN-TAG to a flow
- Able to process a CNM PDU
- Able to back off the transmission rate of the congestion controlled flow based on information contained in the CNM PDU

The reaction point is connected to a destination device (callout 5) by traversing switches A and B (callouts 2 and 3). All egress ports on switches in the path between the reaction point and the destination are configured as congestion points (callouts 4). All congestion points are configured for a CNPV. In our example, CNPV 6 is mapped to transmit queue 6 for all congestion points.

Up to seven 802.1p priorities mapped to a port's transmit queues can be configured as CNPVs. At least one 802.1p priority on a port must be a non-congestion aware priority. Any non-congestion aware priority can be used as a congestion notification alternate priority. When a packet that does not belong to a congestion controlled flow has the same priority as a CNPV configured on a

congestion notification domain edge ingress port, it must be remapped to an alternate priority to defend against a false triggering of a congestion notification by a non-congestion controlled flow.

In [Figure 16-2](#), the reaction point tags the traffic with the CNPV 6 mapped to queue 6 and transmits it to the destination. The congestion point at switch B identifies congestion and creates a CNM PDU packet that is sent back to the reaction point. When the reaction point receives the CNM PDU, it uses the information contained in the CNM PDU to back off the reaction point transmission rate for the queue associated with the congestion controlled flow that triggered the CNM PDU.

A CN-TAG helps the source identify the flow. The CN-TAG is optional for packets transmitted by the source and, if present, may contain a flow-ID. All CNM PDU packets sent from the congestion point back to the reaction point contain a CN-TAG. The transmit queue that detects the congestion will use the flow-ID from the CN-TAG, if it is present, when generating the CNM PDU back to the source. If the congestion controlled flow does not contain a CN-TAG, congestion notification sets the flow-ID to 0.

The traffic monitored by congestion notification must be isolated to its own CNPV. Traffic that does not support congestion notification must not be placed on a CNPV mapped queue. The source of this non-congestion aware priority traffic would not understand the CNM PDU being transmitted back from the congestion point, defeating the purpose of congestion notification. LLDP may be used to auto-determine the port's capability. The information that LLDP provides ensures that non-supported flows are not placed in the congestion aware traffic classes mapped to a CNPV.

Congestion notification advertises its support and state via a TLV type 127 defined in LLDP. If a non-congestion aware priority frame ingresses a congestion notification domain edge port, with the same priority as a configured CNPV on that port, congestion notification will use an alternate priority value to remap the packets to a non-CNPV queue. This alternate priority is configurable for each CNPV.

Using LLDP and alternate priority settings, devices can be configured to remap a priority for all ports away from a congestion notification configured queue and rely upon LLDP to dynamically determine which ports support congestion notification. The protection of the congestion notification domain from non-CN capable port traffic is referred to as congestion notification domain defense. See "[Congestion Notification Domain Defense](#)" on page 16-9 for a detailed discussion of domain defense.

## Implementing Congestion Notification

To implement congestion notification:

- Globally enable congestion notification on the device
- Create congestion notification priority values and optionally specify the creation mode for the CNPVs
- Activate the configured congestion notification priority values on all switch ports
- Optionally modify the transmit priority for the CNM PDU sent by the congestion point to the reaction point
- Optionally modify the default domain defense mode for a priority for the device or on a port basis
- Optionally modify the default priority choice globally or on a port basis
- Optionally modify the default alternate priority for a CNPV globally or on a port basis
- Optionally enable congestion notification LLDP for a CNPV for the device or on a port basis

- If priority choice auto is configured on a port, enable the sending of congestion notification TLVs in the LLDP PDUs for each congestion point port on the system.

## Enabling Congestion Notification

Congestion notification must be globally enabled on the switch for congestion notification to be operational. Use the **set dcb cn global** command to globally enable congestion notification on the switch. This command globally enables congestion notification on the switch, but it does not affect the state of a created CNPV. See “[Congestion Notification Priority Value \(CNPV\)](#)” on page 16-8 for CNPV creation and activation information. Congestion notification is globally enabled by default.

The transmit priority option of the **set dcb cn global** command allows you to modify the transmit priority value of the CNM PDUs sent from the congestion point to the reaction point when congestion is detected. The default value is priority 6.

## Congestion Notification Priority Value (CNPV)

A congestion notification priority value (CNPV) is an 802.1p value configured for congestion notification and mapped to the same queue on all ports that make up the congestion notification domain for that CNPV. There are eight 802.1p values from 0 - 7. The maximum number of CNPVs configurable on a port is seven. There must always be at least one alternate (non-CNPV) priority value per port.

Use the **set dcb cn priority** command, specifying the 802.1p priority to set as the CNPV and optionally specifying a creation mode, to create a CNPV for all ports on the device.

By default, a CNPV is created with auto creation enabled when either the **creation enable** option is specified or no option is specified. “Auto creation enabled” specifies that the port-priority choice mode for each port-priority configured for this CNPV is set to **auto** and the domain defense mode default is set to **edge**. See “[Priority Choice](#)” on page 16-11 for priority choice details. See “[Congestion Notification Domain Defense](#)” on page 16-9 for domain defense mode details.

If **creation disable** is specified when creating the CNPV, the port-priority choice mode for each port-priority is set to **admin** and the domain defense mode default is set to **interior**. Priority choice mode **admin** uses administratively configured settings when determining the domain defense mode for a port.

Use the **set dcb cn priority status** command to activate a CNPV on the switch. Congestion notification will not occur for a CNPV unless the CNPV is activated on the switch. CNPVs are active by default. Priority status defaults to **enable** on creation.

## Alternate Priority

The congestion notification alternate priority is a non-CNPV used to protect the congestion notification domain from a non-congestion controlled flow packet with the same priority as a configured CNPV on the port from triggering congestion notification. At least one 802.1p priority on a port must be a non-CNPV. Any non-CNPV can be used as a congestion notification alternate priority.

When a packet ingresses a port at the edge of a congestion notification domain and has the same priority as a CNPV configured on the ingress port, the packet’s priority must be remapped to an alternate priority. Should a non-congestion notification packet trigger congestion in a CNPV queue, the source for this packet will not know what to do with the CNM PDU it receives back from the congestion point. The remapping of the priority to a non-CNPV value at the congestion notification domain edge guards against this possibility.

An alternate priority can be set both globally or on a port basis. The global alternate priority is only used if the port-priority choice set using the **set dcb cn port-priority choice** command is set

to **default** and the priority choice set using the **set dcb cn priority choice** command is set to **admin**. The port-priority alternate priority is only used if the port-priority choice is set to **admin**. (See “[Priority Choice](#)” on page 16-11 for a priority choice discussion.) This administratively set global or port-based alternate priority defaults to 0. Otherwise, the auto alternate priority is used. The auto alternate priority for a CNPV defaults to either the next lowest non-CNPV priority if a lower one exists on the port or the next highest non-CNPV priority on the port.

Use the **set dcb cn priority alt-pri** command to change the global alternate priority default.

Use the **set dcb cn port-priority alt-pri** command to change the alternate priority default on a port basis. This command overrides any global configuration for the specified port.

## Congestion Notification Domain Defense

A congestion notification domain defense provides a means of defending a congestion notification domain against incoming frames from outside of the domain. Domain defense assumes:

- That every bridge along a path between two congestion aware end-stations, using a particular CNPV, is properly configured for congestion notification and therefore belongs to the congestion notification domain
- That every bridge ensures that frames not configured for a CNPV use different queues than the CNPV configured queues for those devices

Domain defense protects the boundaries of a congestion notification domain by preventing frames not in a congestion controlled flow from entering congestion point controlled queues. Domain defense takes advantage of the ability to change the priority value based upon whether or not the port’s neighbor is also configured with the same CNPV. If a frame with the same priority as the CNPV is not in the congestion controlled flow, the frame priority is changed to the configured alternate priority for that CNPV.

A default domain defense mode is configured at each congestion point port.

There are four possible domain defense modes depending upon whether the CNPV is configured for the congestion point, whether a given congestion point knows the congestion notification state of its neighbor, and where the congestion point port is located in the congestion notification domain:

- **Disabled** – The domain defense mode state on a port for which congestion notification is disabled. The priority is not a CNPV. Congestion notification does not control priority remapping of input frames on this port. CN-TAGs are neither added by an end station nor removed by a bridge. Disabled mode is only set administratively.
- **Edge** – The domain defense mode configured on a congestion point port that resides at the edge of the congestion notification domain. All frames ingressing the edge of a congestion notification domain by definition do not belong to a congestion controlled flow for this domain. On this port for the given CNPV, congestion notification controls priority remapping. The input frame priority parameters are remapped to an alternate (non-CNPV) value. CN-TAGs are not added by an end station, and are removed from frames before being output by a bridge. This mode is optional for an end station.

The end result is that ports configured for edge domain defense protect the congestion notification domain by reassigning the 802.1p priorities of non-CNPV ingressing frames to alternate priorities when they are the same as an edge port CNPV.

When configured for LLDP dynamic congestion notification, a port will be auto configured as edge when the neighbor is not configured for this CNPV.

- **Interior** – The domain defense mode configured on a congestion point port that resides within the congestion notification domain between the flow’s source reaction point and the destination end-station. This port does not yet know whether its neighbor is able to receive a

CN-TAG in frames sent to it. On this port for the given CNPV, the input frame priority parameters are not remapped. CN-TAGs are not added by an end station, and are removed from frames before being output by a bridge.

The Interior defense mode is a transition state within an LLDP dynamic congestion notification configuration. Once the port completes congestion notification negotiation with its neighbor, the defense mode transitions to interior-ready.

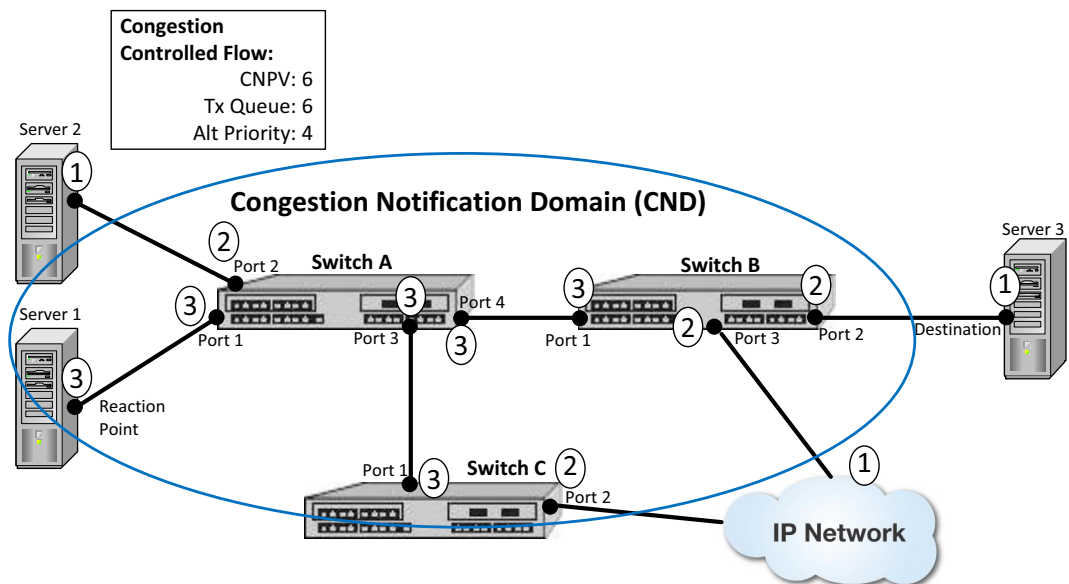
- **Interior-Ready** – The domain defense mode configured on an interior congestion port that knows its neighbor is able to receive a CN-TAG in frames sent to it. On this port for the given CNPV, the input frame priority parameters are not remapped. CN-TAGs can be added by an end station, and are not removed from frames by a bridge.



**Note:** 7100-Series Hardware bonding ports must be manually configured for congestion notification domain defense to interior-ready.

Figure 16-3 on page 16-10 provides a dynamic domain defense mode configuration overview.

**Figure 16-3 Congestion Notification Domain Defense Mode Overview**



1 Non-CN configured port      2 Edge defense port      3 Interior ready defense port

In Figure 16-3, there are two packet flow sources. One of the flow sources is a reaction point configured for CNPV 6 and mapped to queue 6 (Server 1). The second flow source is not configured for congestion notification (Server 2).

There are two paths between the two packet flow sources and the destination. The first path is from the two flow sources to the destination through switches A and B. The second path is from the flow sources to the destination through switches A and C, an IP network cloud, and switch B.

There are three flow discussions that can be derived from Figure 16-3.

- **Server 2 non-congestion notification flow** – The 802.1p priority 6 frame sourced at Server 2 is a non-CN frame because its source is not a reaction point within a congestion notification domain. As the frame enters port 2 Switch A, because port 2 is a domain edge port, and the frame priority agrees with CNPV for this domain, the frame priority (6) is changed to the

alternate priority value (priority 4). The frame transits the remainder of the path to the destination incapable of triggering congestion notification.

- **Server 1 congestion notification flow (Switch A/Switch B path)** – The CNPV 6 frame sourced at Server 1 (reaction point) is a congestion notification frame. As the frame transits to the destination, both ingress ports are configured for the interior-ready defense mode because they have successfully negotiated CNPV 6 with their peers. The CNPV value is not changed to an alternate priority when ingressing interior-ready ports. The frame exits the congestion notification domain for CNPV 6 at port 2, Switch B, and arrives at the destination with its priority unchanged.

Should congestion occur at port 4 of Switch A or port 2 of Switch B, a CNM PDU will be sent back to the reaction point which will back off the flow transmit rate so long as it receives CNM PDUs from the congestion point.

- **Server 1 congestion notification flow (IP Network path)** – The CNPV 6 frame sourced at Server 1 (reaction point) is a congestion notification frame. As the frame transits to the destination, Switch A and Switch C ingress ports are configured for the interior-ready defense mode because they have successfully negotiated CNPV 6 with their peers. The frame leaves the congestion notification domain at port 2, Switch C. The CN-Tag (if present) is stripped and the priority is unchanged. Assuming there are no policy or other reasons why the priority would be changed in the IP Network cloud, the frame arrives at port 3, Switch B with a priority of 6. Because port 3, Switch B is configured for edge defense and the frame priority is the same as the congestion notification domain CNPV 6, the frame priority is changed to the alternate priority 4 no longer capable of triggering congestion notification within this domain. The frame transits to the destination with a priority of 4.

Should congestion occur at port 3, Switch A or port 2, Switch C, a CNM PDU will be sent back to the reaction point which will back off the flow transmit rate so long as it receives CNM PDUs from the congestion point. Should congestion occur at port 2, Switch B, congestion notification is not triggered because the priority is now alternate priority 4.

Port defaults for domain defense are determined by priority choice. See “[Priority Choice](#)” on page 16-11 for details.

Defaults for domain defense can be administratively configured by priority on a port or for all priorities on a port. A default domain defense can be globally configured per CNPV for all ports using the **set dcb cn priority defense** command. A default domain defense can be set on a port basis for all CNPVs on that port using the **set dcb cn port-priority defense** command.

## Priority Choice

There are two priority choice modes that determine how domain defense is configured globally on the switch and three priority choice modes on a port basis. The priority choice mode can be set to:

- **Admin** – Domain defense is administratively configured. When defense choice is set to **admin**, defense mode defaults to **interior** on all ports. Admin can be configured both globally and on a port basis.
- **Auto** – Domain defense is dynamically configured using LLDP. When defense choice is set to **auto**, defense mode defaults to **edge** on all ports. Auto can be configured both globally and on a port basis.
- **Default** – Domain defense is based upon the creation setting (**enable** or **disable**) used when the CNPV is created. If **creation enable** is set, domain defense defaults to **auto**. If **creation disable** is set, domain defense defaults to **admin**.



[Table 16-1](#) cross-references port-priority choice and priority choice settings with the default defense mode and alternate priority settings.

**Table 16-1 Choice, Defense Mode, and Alternate Priority Cross-Reference**

If Port-Priority Choice is:	And Global Priority Choice is:	Then Defense Mode is:	And Alternate Priority is:
auto	auto or admin	auto defense mode (Default: edge)	auto alternate priority
admin	auto or admin	port-priority admin defense mode (Default: disabled)	port-priority alternate priority
default	auto	auto defense mode (Default: edge)	auto alternate priority
default	admin	priority admin defense mode (Default: interior)	priority alternate priority

For example, if the global priority choice is set to auto and the port-priority choice is set to default (row 3 of [Table 16-1](#)), both the defense mode and alternate priority are auto chosen. In this case, the defense mode would default to edge and the alternate priority would default to the next lowest non-CNPV value, or if no lower one exists, the next highest non-CNPV value.

Priority choice on a global basis is configured using the **set dcb cn priority choice** command.

Priority choice on a port basis is configured using the **set dcb cn port-priority choice** command.

## LLDP

LLDP can be used to dynamically enable domain defense on a port for a CNPV. Since Congestion Notification requires every node in the path to be capable of congestion notification support, the nearest bridge LLDP address is used in advertising the support. LLDP TLV type 127 is used by congestion notification and includes a:

- Per-priority CNPV indicator – An 8-bit field where each bit represents the priority for the port and if it's capable of congestion notification.
- Per-priority ready indicator – An 8-bit field where each bit represents the priority for the port and if it's ready for congestion notification (set to 1 if the port is enabled for CNPV).

If a congestion notification capable port receives a TLV type 127 signaling that its neighbor is capable on the same CNPV, the ready indicator is set and the port transitions from the edge to the interior ready domain defense mode for that CNPV.

For dynamic configuration of domain defense to take place you must:

- Assure that the global priority choice is set to auto (default setting when creating a CNPV in creation enable mode)
- Enable congestion notification LLDP on the device using the **set dcb cn priority lldp** command (defaults to enabled)
- Enable the sending of congestion notification TLVs on each congestion point port using the **set lldp port tx-tlv congestion-notif** command (defaults to disabled)

## Congestion Point Queue

The settings for a subset of parameters associated with a congestion point queue can be modified. Congestion point indexes are assigned to congestion point queues which have one or more CNPVs mapped to them on a given port. These congestion queues can be configured using the assigned congestion point index. The congestion point index is one greater than the lowest CNPV which is mapped to the congestion notification queue. For example: if CNPVs 4 and 7 are configured for a transmit queue, the congestion point index is 5. A queue profile which supports the configuring of a minimum sample option can also be associated with the congestion point configuration.

Congestion point queues are configured with the **set dcb congestion-point** command. You can configure the following:

- **Set point** – A target value for the number of octets in the congestion point queue. CNM PDUs are transmitted to the sources of frames queued in this congestion point's queue in order to keep the total number of octets stored in the queue at the set point value.
- **Weight** – An integer value used for calculating  $W$ .  $W$  is the weight to be given to the change in queue length when calculating a measure of transmit queue congestion known as quantized feedback (Fb) as defined in the IEEE 802.1Q-2011 standard.

The weight option is an integer value from which  $W$  is derived.  $W$  is equal to two to the power of the weight value specified here. Thus, if weight equals a -1,  $W = 1/2$ .  $W$  can be between the values specified by Min Weight and Max Weight for this device as displayed by the **show dcb cn q-profile** command. See the IEEE 802.1Q-2011 standard for a detailed discussion for  $W$ , weight, and Fb.

A maximum of two congestion notification queue profiles can be associated with a congestion point queue configuration. The default queue profile on the the 7100-Series has an identifier of **0.1** and can be displayed using the **show dcb cn q-profile** command. This default queue profile can not be modified. A second queue profile, **1.1**, can be configured. The queue profile is configured using the **set dcb cn q-profile** command. See "[Congestion Notification Queue Profile](#)" on page 16-13 for a congestion notification queue profile discussion.

Use the **show dcb cn congestion-point** command to display congestion point queue information.

## Congestion Notification Queue Profile

Each congestion notification queue belongs to one of two possible queue profiles on the 7100-Series. A queue profile is a management object containing congestion notification queue configuration. A queue profile is named based upon an index value and the queue type in dotted notation. The 7100-Series queue type is always 1. For example: the queue profile name for queue profile index 0 would be **0.1**.

The queue profile is configured with the **set dcb cn q-profile** command. You can configure the minimum sample parameter, which specifies the minimum number of octets to enqueue in the congestion point queue between transmissions of CNM PDUs. The default value is 150000 octets.

A queue profile is applied to a congestion queue configuration by specifying its identifier when configuring a congestion point queue using the **set dcb cn congestion-point** command.

Use the **set dcb cn q-profile** command to modify the minimum sample parameter for a group of congestion point queues.

Use the **show dcb cn q-profile** command to display congestion point queue profile information.

## Congestion Notification Configuration Example

The following CLI input shows how to configure each switch for the dynamic domain defense as presented in [Figure 16-3](#) on page 16-10.

```
7100 Chassis(rw)->set dcb cn priority 6
7100 Chassis(rw)->set dcb cn priority 6 status enable
7100 Chassis(rw)->set lldp port tx-tlv congestion-notif *.*.*
```

This example assumes:

- This congestion notification configuration is a default configuration on all switches using LLDP to dynamically configure domain defense on all ports.
- That CoS priority 6 is mapped to transmit queue 6 on all ports within the congestion notification domain. Should there be other CNPV values configured, these may use transmit queue 6, but no non-CNPV priority is mapped to transmit queue 6.



**Note:** Congestion notification and congestion notification LLDP are globally enabled by default. A non-auto alternate priority is not configurable when using auto congestion notification.

## Configuring Data Center Bridging

[Table 16-2](#) lists the 7100-Series device default Data Center Bridging configuration settings.

**Table 16-2 Default Data Center Bridging Configuration Settings**

Parameter	Description	Default Value
Application Priority	Advertises to the LLDP peer a preferred priority for frames carrying application-specific traffic.	disabled on all ports
CN alternate priority	The priority a packet is mapped to if a non-CNPV packet ingresses a port set for domain defense mode edge.	relative to CNPV, either the next lower non-CNPV priority, if it exists, or the next higher.
CN transmit priority	The priority assigned to the CNM PDUs sent by the congestion point back to the reaction point when congestion is detected.	7
CNPV creation status	Specifies whether auto choice is enabled or disabled on the switch for the created congestion notification priority value (CNPV).	enabled
CNPV status	Specifies whether a configured CNPV is activated (enable) or not activated (disable) on the switch.	activated (enable)
congestion notification domain defense mode	Determines whether the ingressing packet priority is remapped away from a configured CNPV value on the port and whether a CN-TAG can be added or removed from the packet.	Edge if port-priority choice is set to auto; Disabled if port-priority choice is set to admin; Determined by the default choice of the priority, if port-priority is set to default.
congestion notification global status	Specifies whether congestion notification is enabled globally on the switch.	enabled

**Table 16-2 Default Data Center Bridging Configuration Settings (continued)**

Parameter	Description	Default Value
congestion notification LLDP	Used to send LLDP CN TLVs for all CNPVs on the port. The actual sending of the TLVs is disabled by default.	enabled
congestion queue minimum sample	Specifies the minimum number of octets to enqueue in the congestion point queue between transmissions of CNM PDUs	150000 octets
Enhanced Transmission Selection (ETS)	Provides for the designation of two or more traffic class queues to be allocated for bandwidth that will not be serviced until all non-ETS queues are empty.	No ETS class queues are configured
priority choice	Specifies the method for selecting domain defense on all device ports.	auto
Priority-based Flow Control (PFC)	Provides for flow control per hardware egress queue that the PFC priority is mapped to.	disabled on all ports

Table 16-3 lists Data Center Bridging configuration commands.

**Table 16-3 Data Center Bridging Configuration**

Task	Command
Optionally, map traffic classes to ETS groups in the CoS transmit queue port group configuration and specify the bandwidth allocation for the group.	<b>set cos port-config txq</b> <i>group-type-index</i> [ <b>name</b> <i>name</i> ] [ <b>enhanced-groups</b> <i>group-id</i> ] [ <b>enhanced-percentage</b> <i>bandwidth</i> ]
Optionally, enable priority-based flow control for a specified port and priority or modify the link delay allowance.  It is not recommended that you modify the link delay allowance.	<b>set dcb pfc</b> <i>port-string</i> <b>priority enable</b>   <b>link-delay-allowance</b> <i>bits</i>
Optionally, set the priority to be advertised to the peer of the specified port for the specified application.	<b>set dcb appPri</b> <i>port-string</i> <b>protocol</b> { <i>ethertype</i>   <b>tcp</b>   <b>udp</b>   <b>l4port</b> } <b>protocol-id</b> <i>protocol-id</i> <b>priority</b> <i>priority</i>
Set Data Center Bridging optional LLDP-DCB TLVs to be transmitted in LLDP PDUs by the specified port or ports.	<b>set lldp port tx-tlv</b> {[ <b>enhanced-trans-config</b> ] [ <b>enhanced-trans-rec</b> ] [ <b>application-pri</b> ] [ <b>priority-flowctrl</b> ] [ <b>congestion-notif</b> ]} <i>port-string</i>

Table 16-4 lists Congestion Notification global configuration commands. All Congestion Notification commands can be entered from any command mode.

**Table 16-4 Congestion Notification Global Configuration**

Task	Command
Globally enable congestion notification on the switch.	<b>set dcb cn global enable</b>
Optionally, modify the priority for Congestion Notification Messages sent from the congestion point back to the congestion notification reaction point.	<b>set dcb cn global tx-priority</b> <i>tx-priority</i>

**Table 16-4 Congestion Notification Global Configuration (continued)**

Task	Command
Optionally, modify congestion point queue parameters for queues associated with a congestion point queue configuration or associate a congestion notification queue profile with the configuration.	<b>set dcb cn congestion-point</b> <i>port-string cp-index</i> [ <b>set-point</b> <i>set-point</i> ] [ <b>weight</b> <i>weight</i> ] [ <b>qp-index</b> <i>qp-index</i> ]
Optionally, modify congestion notification queue profile parameters associated with up to two queue profiles on the 7100-Series device	<b>set dcb cn q-profile</b> <i>qp-identifier</i> [ <b>min-sample</b> <i>min-sample</i> ]

Table 16-4 lists Congestion Notification global priority configuration commands. All Congestion Notification commands can be entered from any command mode.

**Table 16-5 Congestion Notification Global Priority Configuration**

Task	Command
Configure up to seven 802.1p priorities as Congestion Notification Priority Value (CNPV) on the switch. Optionally, configure priority choice mode to <b>admin</b> using the creation <b>disable</b> option. Priority choice defaults to <b>auto</b> .	<b>set dcb cn priority</b> <i>priority</i> [ <b>creation</b> { <b>enable</b>   <b>disable</b> }]
Activate configured CNPVs for all ports on the switch.	<b>set dcb cn priority</b> <i>priority</i> <b>status enable</b>
Optionally modify the priority choice mode on all device ports for the specified CNPV. Priority choice globally defaults to <b>auto</b> .	<b>set dcb cn priority</b> <i>priority</i> <b>choice</b> { <b>admin</b>   <b>auto</b> }
Optionally, modify the global default congestion notification alternate priority for the specified CNPV.	<b>set dcb cn priority</b> <i>cnpv alt-priority alt-priority</i>
Optionally, administratively configure a global default domain defense mode for all ports on the switch for the specified priority.	<b>set dcb cn priority</b> <i>priority</i> <b>defense</b> { <b>disabled</b>   <b>edge</b>   <b>interior</b>   <b>interior-ready</b> }
If using the default priority choice of <b>auto</b> , enable auto configuration of domain defense operation using LLDP.	<b>set dcb cn priority</b> <i>priority</i> <b>lldp enable</b>

Table 16-6 lists Congestion Notification port priority configuration commands.

**Table 16-6 Congestion Notification Port Priority Configuration**

Task	Command
Optionally, for all priorities, modify the default method for determining how the domain defense is selected for the specified port.	<b>set dcb cn port-priority</b> <i>port-string</i> <b>choice</b> { <b>admin</b>   <b>auto</b>   <b>default</b> }
Optionally modify the alternate priority for the specified port and CNPV.	<b>set dcb cn port-priority</b> <i>port-string priority alt-priority alt-priority</i>
Optionally, administratively modify the default domain defense mode for the specified priority on the specified port.	<b>set dcb cn port-priority</b> <i>port-string priority</i> <b>defense</b> { <b>disabled</b>   <b>interior</b>   <b>interior-ready</b>   <b>edge</b> }
If using the default priority choice of <b>auto</b> , enable auto configuration of domain defense operation using LLDP for a specific priority and port.	<b>set dcb cn port-priority</b> <i>port-string priority</i> <b>lldp</b> { <b>enable</b>   <b>disable</b> }

**Table 16-6 Congestion Notification Port Priority Configuration (continued)**

Task	Command
Optionally, set the method that determines how the domain defense is selected for the specified priority on the specified port.	<b>set dcb cn port-priority</b> <i>port-string</i> <b>priority</b> <i>priority choice</i> {admin   auto   default}

[Table 16-7](#) lists Data Center Bridging display commands.

**Table 16-7 Data Center Bridging Display Commands**

Task	Command
Display the ETS group CoS transmit queue port group mappings by the port group and type.	<b>show cos port-config txq</b> <i>group-type-index</i>
Display priority-based flow control table entries.	<b>show dcb pfc</b> [ <i>port-string</i> ][ <b>-interesting</b> ] [ <b>link-delay-allowance</b> ]
Display Application Priority table entries by port.	<b>show dcb appPri</b> <i>port-string</i>
Display LLDP port Data Center Bridging, priority-based flow control, Application Priority, or congestion notification transmit TLV support.	<b>show lldp port tx-tlv</b> [ <b>data-center-bridging</b> ] [ <b>priority-flowctrl</b> ] [ <b>application-pri</b> ] [ <b>congestion-notif</b> ]
Display the local or remote system information, including ETS information, stored for one or more ports.	<b>show lldp port</b> { <b>local-info</b>   <b>remote-info</b> ] [ <i>port-string</i> ]
Display the global status of congestion notification on the switch.	<b>show dcb cn global</b>
Display the configuration and status of congestion notification priority values on the switch.	<b>show dcb cn priority</b> [ <i>priority</i> ] [ <b>-interesting</b> ]
Display congestion notification port level defense mode configuration.	<b>show dcb cn port-priority</b> [ <i>port-string</i> [ <i>priority</i> ]] [ <b>priority</b> <i>priority</i> ] [ <b>-interesting</b> ]
Display the configuration and status of congestion notification congestion points.	<b>show dcb cn congestion-point</b> [ <i>port-string</i> [ <i>cp-id</i> ]] [ <b>stats</b> ]
Display the congestion notification queue profile configuration stats.	<b>show dcb cn q-profile</b> [ <i>profile-id</i> ]

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

## Terms and Definitions

[Table 16-8](#) lists terms and definitions used in this DCB configuration discussion.

**Table 16-8 Data Center Bridging (DCB) Configuration Terms and Definitions**

Term	Definition
Application Priority	A DCB feature that provides for the advertisement to the peer of a preferred priority to be applied to frames carrying application-specific traffic. The peer must support the LLDP willing bit.
Congestion Notification (CN)	A DCB feature, as defined in IEEE 802.1Q-2011, that allows a device to detect congestion at a switch congestion point and transmit a Congestion Notification Message back to the reaction point indicating the reaction point should back off the traffic for that flow.

**Table 16-8 Data Center Bridging (DCB) Configuration Terms and Definitions (continued)**

Term	Definition
Congestion Notification Message (CNM)	A PDU message sent from the congestion point back to the reaction point to back off on a transmitting the flow when congestion is detected in a congestion notification enabled context.
congestion point	An egress transmit point configured for congestion notification.
Data Center Bridging (DCB)	A group of features that enhance Ethernet technology by enabling the convergence of various applications in data centers, such as Local Area Networks (LAN), Storage Area Networks (SAN), and advanced application High Performance Computing (HPC) onto a single interconnect technology, by providing enhancements to existing 802.1 bridge specifications.
Data Center Bridging Exchange (DCBX)	A protocol that allows Ethernet devices to detect DCB capability on the peer device, as well as DCB configuration between peer devices.
Enhanced Transmission Selection (ETS)	A DCB feature that provides a common management framework for assignment of bandwidth to 802.1p CoS-based traffic classes (IEEE 802.1Qaz).
LLDP willing bit	An LLDP attribute that when enabled instructs the local device to use the peer Application Priority settings contained in received TLVs. The LLDP willing bit is not currently supported on Extreme Networks switches.
Priority-based Flow Control (PFC)	A DCB feature that provides a link level flow control mechanism that can be controlled independently for each Class of Service (CoS), with a goal of ensuring zero loss under congestion in DCB networks (IEEE 802.1Qbb).
priority or port-priority choice	The method used to select the domain defense either globally (priority) or on a port-basis (port-priority). Both global and port-based choice supports auto or administrative selection of the domain defense. On a port-basis, domain defense selection can be set to the default creation value set when the CNPV was created.
reaction point	The flow source device capable of optionally adding a CN-TAG to the flow, has the ability to process a CNM PDU, and is able to throttle its transmission rates based on information contained in the CNM PDU. .

## Simple Network Management Protocol (SNMP) Configuration

This chapter provides information about configuring and monitoring SNMP on Extreme Networks 7100-Series devices.

For information about...	Refer to page...
<a href="#">Using SNMP in Your Network</a>	17-1
<a href="#">SNMP Concepts</a>	17-2
<a href="#">SNMP Support on 7100-Series Devices</a>	17-4
<a href="#">Configuring SNMP</a>	17-7
<a href="#">Reviewing SNMP Settings</a>	17-20

### Using SNMP in Your Network

The Simple Network Management Protocol (SNMP) is an application layer protocol that facilitates the exchange of management information between network devices. The most widely used management protocol on Internet Protocol (IP) networks, it helps you monitor network performance, troubleshoot problems, and plan for network growth.

SNMP's simplicity lies in the fact that it uses a basic set of command messages to relay notifications of events and error conditions over a connectionless communication link.

Most network devices support the three versions of the protocol: SNMPv1, SNMPv2c, and SNMPv3. The latest version, SNMPv3, provides enhanced security and administrative features as described in this document.

SNMP is a simple, cost-effective tool for monitoring your network devices for conditions that warrant administrative attention. It is widely used because it is:

- Easily integrated into your existing LAN topology
- Based on an open standard, making it non-proprietary and well documented
- Flexible enough to communicate the specific conditions you need monitored in your network
- A common management platform supported by many network devices



## High-Level Configuration Process

You can implement SNMP on Extreme Networks switching devices using simple CLI commands as described in this chapter. The configuration process involves the following tasks:

1. Creating users and groups allowed to manage the network through SNMP
2. Setting security access rights
3. Setting SNMP Management Information Base (MIB) view attributes
4. Setting target parameters to control the formatting of SNMP notification messages
5. Setting target addresses to control where SNMP notifications are sent
6. Setting SNMP notification parameters (filters)
7. Reviewing SNMP statistics

## SNMP Concepts

It is helpful to understand the following SNMP concepts:

For information about...	Refer to page...
<a href="#">Manager/Agent Model Components</a>	17-2
<a href="#">Message Functions</a>	17-2
<a href="#">Access to MIB Objects</a>	17-3

### Manager/Agent Model Components

SNMP provides a message format for communication between managers and agents, which use a MIB and a relatively small set of commands to exchange information. The SNMP manager can be part of a network management system, such as Extreme Networks NetSight, while the agent and MIB reside on the switch.

The SNMP agent acts upon requests from the manager to either collect data from the MIB or to set data into the MIB. A repository for information about device parameters and network data, the MIB is organized in a tree structure in which individual variables are represented as leaves on the branches. A unique object identifier (OID) distinguishes each variable in the MIB and is the means by which the manager and agent specify which managed elements are changed.

An agent can send unsolicited notification messages (also known as traps or informs) alerting the SNMP manager to a condition on the network. These conditions include such things as improper user authentication, restarts, link status (up or down), MAC address tracking, closing of a TCP connection, loss of connection to a neighbor, or other significant events.

### Message Functions

SNMP uses five basic message types (Get, Get Next, Get Response, Set, and Trap) to communicate between the manager and the agent. The Get and Get Next messages allow the manager to request information for a specific variable. The agent, upon receiving a Get or Get Next message, will issue a Get Response message to the manager with either the information requested or an error indication about why the request cannot be processed.

A Set message allows the manager to request a change to a specific variable. The agent then responds with a Get Response message indicating the change has been made or an error indication about why the change cannot be made.

A trap or inform message allows the agent to spontaneously inform the manager of an “important” event in the network.

The SNMP manager and agent use information in the MIB to perform the operations described in [Table 17-1](#).

**Table 17-1 SNMP Message Functions**

Operation	Function
get-request	Retrieves a value from a specific variable.
get-next-request	Retrieves a value from a variable within a table. <sup>1</sup>
get-bulk-request <sup>2</sup>	Retrieves large blocks of data, such as multiple rows in a table, that would otherwise require the transmission of many small blocks of data.
get-response	Replies to a get-request, get-next-request, and set-request sent by a management station.
set-request	Stores a value in a specific variable.
trap   inform <sup>3</sup>	Unsolicited message sent by an SNMP agent to an SNMP manager when an event has occurred.

1. With this operation, an SNMP manager does not need to know the exact variable name. A sequential search is performed to find the needed variable from within a table.

2. The get-bulk operation is only supported in SNMPv2c or later.

3. Inform notifications are only supported in SNMPv3.

## Trap Versus Inform Messages

As compared to earlier versions, SNMPv3 provides a higher degree of reliability for notifying management stations when critical events occur. Traditionally, SNMP agents communicated events to SNMP managers via “traps.” However, if a temporary network problem prevented the manager from receiving the trap, then the trap would be lost. SNMPv3 provides “informs”, which are a more reliable form of traps. The SNMP agent initiates the inform process by sending an inform request to the manager. The manager responds to the inform request to acknowledge receipt of the message. If the inform is not received by the manager, the inform request will timeout and a new inform request will be sent. Subsequent inform requests will be sent as previous requests time-out until either an acknowledgement is received from the manager, or until a pre-specified retry-count is reached.

## Access to MIB Objects

SNMP uses the following authentication methods to grant user access to MIB objects and functions.

### Community Name Strings

Earlier SNMP versions (v1 and v2c) rely on community name strings for authentication. In order for the network management station (NMS) to access the switch, the community string definitions on the NMS must match at least one of the three community string definitions on the switch. A community string can have one of these attributes:

- Read-only (**ro**)—Gives read access to authorized management stations to all objects in the MIB except the community strings, but does not allow write access.

- Read-write (**rw**)—Gives read and write access to authorized management stations to all objects in the MIB, but does not allow access to the community strings.

## User-Based

SNMPv3 provides a User-Based Security Model (USM) which relies on a user name match for authenticated access to network management components.

Refer to “[Security Models and Levels](#)” on page 17-6 for more information.

## SNMP Support on 7100-Series Devices

By default, SNMP Version 1 (SNMPv1) is configured on Extreme Networks switches. The default configuration includes a single community name - public - which grants read-write access to the whole MIB tree for both SNMPv1 and SNMPv2c.

This section provides the following information about SNMP support on Extreme Networks devices:

For information about...	Refer to page...
<a href="#">Versions Supported</a>	17-4
<a href="#">Terms and Definitions</a>	17-5
<a href="#">Security Models and Levels</a>	17-6
<a href="#">Access Control</a>	17-7

## Versions Supported

Extreme Networks devices support three versions of SNMP:

- Version 1 (SNMPv1) — This is the initial implementation of SNMP. Refer to RFC 1157 for a full description of functionality.
- Version 2 (SNMPv2c) — The second release of SNMP, described in RFC 1907, has additions and enhancements to data types, counter size, and protocol operations.
- Version 3 (SNMPv3) — This is the most recent version of SNMP, and includes significant enhancements to administration and security. The major difference between SNMPv3 and earlier versions is that v3 provides a User-Based Security Model (USM) to associate users with managed access to security information. In addition to better security and better access control, SNMPv3 also provides a higher degree of reliability for notifying management stations when critical events occur.

SNMPv3 is fully described in RFC 2571, RFC 2572, RFC 2573, RFC 2574, and RFC 2575.

## SNMPv1 and v2c Network Management Components

The Extreme Networks implementation of SNMPv1 and v2c network management components fall into the following three categories:

- Managed devices (such as a switch).
- SNMP agents and MIBs, including SNMP traps, community strings, and Remote Monitoring (RMON) MIBs, which run on managed devices.

- SNMP network management applications, such as the Extreme Networks NetSight application, which communicate with agents to get statistics and alerts from the managed devices.

## SNMPv3 User-Based Security Model (USM) Enhancements

SNMPv3 adds to v1 and v2c components by providing secure access to devices by authenticating and encrypting frames over the network. The Extreme Networks supported advanced security features provided in SNMPv3's User-Based Security Model are:

- Message integrity — Collects data securely without being tampered with or corrupted.
- Authentication — Determines the message is from a valid source.
- Encryption — Scrambles the contents of a frame to prevent it from being seen by an unauthorized source.

Unlike SNMPv1 and SNMPv2c, in SNMPv3, the concept of SNMP agents and SNMP managers no longer apply. These concepts have been combined into an SNMP entity. An SNMP entity consists of an SNMP engine and SNMP applications. An SNMP engine consists of the following four components:

- Dispatcher — Sends and receives messages.
- Message processing subsystem — Accepts outgoing PDUs from the dispatcher and prepares them for transmission by wrapping them in a message header and returning them to the dispatcher. Also accepts incoming messages from the dispatcher, processes each message header, and returns the enclosed PDU to the dispatcher.
- Security subsystem — Authenticates and encrypts messages.
- Access control subsystem — This component determines which users and which operations are allowed access to managed objects.

## Terms and Definitions

Table 17-2 lists common SNMP terms and defines their use on Extreme Networks devices.

**Table 17-2 SNMP Terms and Definitions**

Term	Definition
community	A name string used to authenticate SNMPv1 and v2c users.
context	A subset of MIB information to which associated users have access rights.
engine ID	A value used by both the SNMPv3 sender and receiver to propagate inform notifications.
group	A collection of SNMP users who share the same access privileges.
inform	A notification message sent by an SNMPv3 agent to a network management station, a console, or a terminal to indicate the occurrence of a significant event, such as when a port or device goes up or down, when there are authentication failures, and when power supply errors occur.
MIB	Management Information Base, a repository for information about device parameters and network data organized in a tree structure.
notify profile	Associates target parameters to an SNMP notify filter to determine who should not receive SNMP notifications. This is useful for fine-tuning the amount of SNMP traffic generated.

**Table 17-2 SNMP Terms and Definitions (continued)**

Term	Definition
OID	Object Identifier, a unique ID distinguishing each variable in the MIB and is the means by which the SNMP manager and agent specify which managed elements are changed.
security level	The permitted level of security within a security model. The three levels of SNMP security are: <ul style="list-style-type: none"> <li>• no authentication required (NoAuthNoPriv)</li> <li>• authentication required (AuthNoPriv)</li> <li>• privacy (authPriv)</li> </ul>
security model	An authentication strategy that is set up for an SNMP user and the group in which the user resides. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP frame.
storage type	Specifies whether an SNMP user entry will be stored in volatile or nonvolatile memory.
taglist	A list of SNMP notify values that link a target (management station IP) address to specific SNMP notifications.
target address	A unique identifier and a specific IP address that will receive SNMP notification messages.
target parameters	A named set of security/authentication criteria used to generate a message to a target.
trap	A notification message sent by an SNMPv1 or v2c agent to a network management station, a console, or a terminal to indicate the occurrence of a significant event, such as when a port or device goes up or down, when there are authentication failures, and when power supply errors occur.
user	A person registered in SNMPv3 to access management information. In v1 and v2c, a user is set with the community name string.
USM	User-Based Security Model, the SNMPv3 authentication model which relies on a user name match for access to network management components.
VACM	View-based Access Control Model, which determines remote access to SNMP managed objects, allowing subsets of management information to be organized into user views.
view	Specifies permission for accessing SNMP MIB objects granted to a particular SNMP user group. View types and associated access rights are: <ul style="list-style-type: none"> <li>• read - view-only access</li> <li>• write - allowed to configure MIB agent contents</li> <li>• notify - send trap messages</li> </ul>

## Security Models and Levels

An SNMP security model is an authentication strategy that is set up for a user and the group in which the user resides. A security level is the permitted level of security within a security model. The three levels of SNMP security on Extreme Networks devices are:

- No authentication required (NoAuthNoPriv)
- Authentication required (AuthNoPriv)
- Privacy (authPriv)

A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP frame. [Table 17-3](#) identifies the levels of SNMP security available on Extreme Networks devices and authentication required within each model.

**Table 17-3 SNMP Security Models and Levels**

Model	Security Level	Authentication	Encryption	How It Works
v1	NoAuthNoPriv	Community string	None	Uses a community string match for authentication.
v2c	NoAuthNoPriv	Community string	None	Uses a community string match for authentication.
v3 / USM	NoAuthNoPriv	User name	None	Uses a user name match for authentication.
	AuthNoPriv	MD5 or SHA	None	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms.
	authPriv	MD5 or SHA	DES	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms. Provides DES 56-bit encryption in addition to authentication based on the CBC-DES (DES-56) standard.

## Access Control

In addition to the [Security Models and Levels](#) described above, the Extreme Networks implementation of SNMP also provides a View-based Access Control Model (VACM), which determines remote access to managed objects. VACM allows you to organize subsets of management information into “views.” Management information that is in a user’s view gives the user the corresponding access level to that management information: either read, write, or notify. Individual users can be organized into groups for whom you can pre-define what views are available based on the security model and security level used to request access. In this way, VACM allows you to permit or deny access to any individual item of management information depending on a user’s group membership and the level of security provided by the communications channel.

## Configuring SNMP

This section provides the following information about configuring SNMP on Extreme Networks devices:

For information about...	Refer to page...
<a href="#">Configuration Basics</a>	<a href="#">17-8</a>
<a href="#">How SNMP Processes a Notification Configuration</a>	<a href="#">17-8</a>
<a href="#">SNMP Defaults</a>	<a href="#">17-9</a>
<a href="#">Configuring SNMPv1/SNMPv2c</a>	<a href="#">17-9</a>
<a href="#">Configuring SNMPv3</a>	<a href="#">17-11</a>
<a href="#">Configuring Secure SNMP Community Names</a>	<a href="#">17-18</a>

## Configuration Basics

Completing an SNMP configuration on an Extreme Networks device involves defining users who will be authorized to receive SNMP notifications about network events, associating security (target) parameters, access rights and MIB views to those users, and specifying an IP address where they will receive notifications. The basic steps in this process are:

1. Creating a name that will act as an SNMP user password:
  - This will be a **community** name for an SNMPv1 or v2c configuration, or
  - A **user** name for an SNMPv3 configuration.
2. Creating a group for the user named in [Step 1](#).
3. Creating access rights for the user group named in [Step 2](#).
4. Defining MIB view(s) for the user group.
5. Creating a target parameters entry to associate security and authorization criteria to the users created in [Step 1](#).
6. Verifying if any applicable SNMP notification entries exist, or creating a new one. You will use this entry to send SNMP notification messages to the appropriate targets configured in [Step 5](#).
7. Creating a target address entry to bind a management IP address to:
  - The notification entry and tag name created in [Step 6](#), and
  - The target parameters entry created in [Step 5](#).



**Note:** Commands for configuring SNMP on Extreme Networks devices are independent during the SNMP setup process. For instance, target parameters can be specified when setting up optional notification filters — even though these parameters have not yet been created with the **set snmp targetparams** command. The steps in this section are a guideline to configuring SNMP and do not necessarily need to be executed in this order.

## How SNMP Processes a Notification Configuration

In order to send a trap or inform notification requested by a MIB code, the SNMP agent requires the equivalent of a trap “door”, a “key” to unlock the door, and a “procedure” for crossing the doorstep. To determine if all these elements are in place, the SNMP agent processes a device configuration as follows:

1. Determines if the “keys” for trap “doors” do exist. The key that SNMP is looking for is the notification entry created with the **set snmp notify** command.
2. Searches for the doors matching such a key and verifies that the door is available. If so, this door is tagged or bound to the notification entry. It was built using the **set snmp targetaddr** command, which specifies the management station IP address to which this door leads, and the “procedure” (**targetparams**) to cross the doorstep
3. Verifies that the description of how to step through the door is, in fact, there. The agent checks **targetparams** entries and determines this description was made with the **set snmp targetparams** command, which tells exactly which SNMP protocol to use and what community or user name to provide.
4. Verifies that the specified name, configured using either the **set snmp community** or **set snmp user** command is available.
5. Sends the notification message to the target address.

## SNMP Defaults

### Device Start Up Configuration

By default, SNMPv1 is configured on Extreme Networks switches. [Table 17-4](#) lists the default configuration parameters, which include a single community name - public - granting read-write access to the whole MIB tree for both SNMPv1 and SNMPv2c.

**Table 17-4 Default Extreme Networks SNMP Configuration**

Parameter	Default Value
Community name	public
Group access privileges	rw (read-write)
Group user name	public
Security model	v1
Security access rights	all (for read, write, and notify access)
MIB view	all (entire MIB tree)

You can revise this default configuration by following the steps described in [“Adding to or Modifying the Default Configuration”](#) on page 17-10.

To take advantage of the advanced security and other features available in SNMPv3, it is recommended that you add to the Extreme Networks default configuration by configuring SNMPv3 as described in [“Configuring SNMPv3”](#) on page 17-11.

Refer also to [“Configuring Secure SNMP Community Names”](#) on page 17-18 for a description of a recommended configuration that will prevent unsecured access to SNMP information.

## Configuring SNMPv1/SNMPv2c

### Creating a New Configuration

[Procedure 17-1](#) shows how to create a new SNMPv1 or SNMPv2c configuration. This example assumes that you haven't any preconfigured community names or access rights.



**Note:** The **v1** parameter in this example can be replaced with **v2** for SNMPv2c configuration.

**Procedure 17-1 New SNMPv1/v2c Configuration**

Step	Task	Command(s)
1.	Create a community name.	<b>set snmp community</b> <i>community</i> [ <b>securityname</b> <i>securityname</i> ] [ <b>context</b> <i>context</i> ] [ <b>transport</b> <i>transport</i> ] [ <b>volatile</b>   <b>nonvolatile</b> ]
2.	Create a security model (VACM) group using the <i>community name</i> you assigned in step 1.	<b>set snmp group</b> <i>groupname</i> <b>user</b> <i>communityname</i> <b>security-model</b> <b>v1</b>
3.	Set security access rights for the VACM group.	<b>set snmp access</b> <i>groupname</i> <b>security-model</b> <b>v1</b> <b>read</b> <i>viewname</i> <b>write</b> <i>viewname</i> <b>notify</b> <i>viewname</i>
4.	Set MIB view attributes.	<b>set snmp view</b> <i>viewname</i> <i>viewname</i> <b>subtree</b> <i>subtree</i>



**Procedure 17-1 New SNMPv1/v2c Configuration (continued)**

Step	Task	Command(s)
5.	Specify the target parameters for SNMP notification message generation.	<b>set snmp targetparams</b> <i>paramset_name</i> <b>user</b> <i>community name</i> <b>security-model v1 message processing v1</b>
6.	Specify the target address to which SNMP notification messages generated using the specified target parameters will be sent.	<b>set snmp targetaddr</b> <i>targetaddr_name</i> <i>ipaddr</i> <b>param</b> <i>paramset_name</i> <b>taglist</b> <i>taglist</i>
7.	Specify a name for this notification entry and bind it to the target address.	<b>set snmp notify</b> <i>notify</i> <b>tag</b> <i>taglist</i>

**Example**

The following example displays an 7100-Series device configuration using the steps in [Procedure 17-1](#). It shows how to:

- Create the community name **public**.
- Assign the **public** user to the group named **groupRW** and the SNMPv1 security model.
- Specify that, if SNMP messages are received with the **public** name string, the view **RW for** read requests, write requests, and notify requests will be applied to this user.
- For the view **RW**, include the MIB subtree denoted with OID **1** and **0.0**, and exclude view access to subtree denoted with OID **1.3.6.1.6.3.13.1** (which is the notification MIB).
- Assign a target parameters entry, **TVv1public**, for security level processing to the **public** community name.
- Create a target address entry named **TVTrap** at IP address **10.42.1.10**, which will use security and authorization criteria contained in the target parameters entry called **TVv1public**, and bind these parameters together with a tag entry called **TVTrapTag**.

```
7100 Chassis(su)->set snmp community public
7100 Chassis(su)->set snmp group groupRW user public security model v1
7100 Chassis(su)->set snmp access groupRW security-model v1 read RW write RW
notify RW
7100 Chassis(su)->set snmp view viewname RW subtree 1
7100 Chassis(su)->set snmp view viewname RW subtree 0.0
7100 Chassis(su)->set snmp view viewname RW subtree 1.3.6.1.6.3.13.1 excluded
7100 Chassis(su)->set snmp targetparams TVv1public user public security-model v1
message processing v1
7100 Chassis(su)->set snmp targetaddr TVTrap 10.42.1.10 param TVv1public taglist
TVTraptag
7100 Chassis(su)->set snmp notify TVTrap tag TVTrapTag
```

**Adding to or Modifying the Default Configuration**

By default, SNMPv1 is configured on Extreme Networks switches. A single community name - **public** - is configured, which grants read-write access to the whole MIB tree for both SNMPv1 and SNMPv2c.

The beginning command sequence in the default configuration is similar to the first part of the previous example. It looks like this:

```
7100 Chassis(su)->set snmp community public
7100 Chassis(su)->set snmp group groupRW user public security-model v1
```

```
7100 Chassis(su)->set snmp access groupRW security-model v1 read All write All
notify All
7100 Chassis(su)->set snmp view viewname All subtree 1
```



**Note:** Any use of the parameter 'All' must be exactly as shown in this example. Any other variation (including, but not limited to, values such as 'all' or 'ALL') will not be valid.

You can modify this default configuration as shown in the following examples.

### Adding a New Community Name

Use these commands to add a new SNMPv1 community name called **newname with the same** permissions as the default configuration:

```
7100 Chassis(su)->set snmp community newname
7100 Chassis(su)->set snmp group groupRW user newname security-model v1
```

Use this command to remove the **public** community name from the default configuration:

```
7100 Chassis(su)->clear snmp community public
```



**Note:** You can leave the **set snmp group groupRW user public security-model v1** statement in the default configuration in case you want to re-activate the **public** community name at some point, or can clear it as well.

Refer to “[Configuring Secure SNMP Community Names](#)” on page 17-18 for a description of a recommended configuration that will prevent unsecured access to SNMP information.

## Configuring SNMPv3

[Procedure 17-2](#) shows how to complete a basic SNMPv3 configuration.

### Procedure 17-2 SNMPv3 Configuration

Step	Task	Command(s)
1.	Create an SNMPv3 user and specify authentication, encryption, and security credentials. <ul style="list-style-type: none"> <li>If <b>remote</b> is not specified, the user will be registered for the local SNMP engine.</li> <li>If <b>authentication</b> is not specified, no authentication will be applied.</li> <li>If <b>privacy</b> is not specified, no encryption will be applied.</li> </ul>	<b>set snmp user</b> <i>user</i> [ <b>remote</b> <i>remoteid</i> ] [ <b>authentication</b> { <i>md5</i>   <i>sha</i> }] [ <i>authpassword</i> ] [ <b>privacy</b> <i>privpassword</i> ]
2.	Create a user group and add the user created in Step 1. <ul style="list-style-type: none"> <li>If storage type is not specified, <b>nonvolatile</b> will be applied.</li> </ul>	<b>set snmp group</b> <i>groupname</i> <b>user</b> <i>user</i> <b>security-model</b> <i>usm</i> [ <i>volatile</i>   <i>nonvolatile</i> ]

**Procedure 17-2 SNMPv3 Configuration (continued)**

Step	Task	Command(s)
3.	<p>Set security access rights for the group.</p> <ul style="list-style-type: none"> <li>If security level is not specified, no authentication will be applied.</li> <li>Only one context, the “default context”, is supported in this release. There is no need to configure this parameter.</li> <li>If <b>read</b> view is not specified none will be applied.</li> <li>If <b>write</b> view is not specified, none will be applied.</li> <li>If <b>notify</b> view is not specified, none will be applied.</li> <li>If storage type is not specified, entries will be stored as permanent and will be held through device reboot.</li> </ul>	<pre>set snmp access <i>groupname</i> security-model usm [noauthentication   authentication   privacy] [exact   prefix] [read <i>readviewname</i>] [write <i>writeviewname</i>] [notify <i>notifyviewname</i>] [volatile   nonvolatile]</pre>
4.	<p>Define views created in Step 3.</p> <ul style="list-style-type: none"> <li>If not specified, <b>mask</b> will be set to empty.</li> <li>If not specified, subtree use will be <b>included</b>.</li> <li>If storage type is not specified, <b>nonvolatile</b> (permanent) will be applied.</li> </ul>	<pre>set snmp view <i>viewname</i> <i>viewname</i> subtree subtree [mask <i>mask</i>] [included   excluded] [volatile   nonvolatile]</pre>
5.	<p>Set SNMP target parameters.</p> <ul style="list-style-type: none"> <li>If not specified, security level will be set to <b>noauthentication</b>.</li> <li>If not specified, storage type will be set to <b>nonvolatile</b>.</li> </ul>	<pre>set snmp targetparams <i>paramset_name</i> user user security-model usm message-processing v3 [noauthentication   authentication   privacy] [volatile   nonvolatile]</pre>
6.	<p>Set the SNMP target address for notification message generation.</p> <ul style="list-style-type: none"> <li>If not specified, <i>udpport</i> will be set to <b>162</b>.</li> <li>If not specified, <i>mask</i> will be set to <b>255.255.255.255</b>.</li> <li>If not specified, <i>timeout</i> will be set to <b>1500</b> (15 seconds).</li> <li>If not specified, number of <i>retries</i> will be set to <b>3</b>.</li> <li>If <b>taglist</b> is not specified, none will be set.</li> <li>If not specified, storage type will be <b>nonvolatile</b>.</li> </ul>	<pre>set snmp targetaddr <i>targetaddr_name</i> <i>ipaddr</i> param <i>paramset_name</i> [udpport <i>udpport</i>] [mask <i>mask</i>] [timeout <i>timeout</i>] [retries <i>retries</i>] [taglist <i>taglist</i>] [volatile   nonvolatile]</pre>
7.	<p>Set SNMP notification parameters.</p> <ul style="list-style-type: none"> <li>If not specified, message type will be set to <b>trap</b>.</li> <li>If not specified, storage type will be set to <b>nonvolatile</b>.</li> </ul>	<pre>set snmp notify <i>notify</i> tag <i>tag</i> [trap   inform] [volatile   nonvolatile]</pre>

The following example is an 7100-Series device configuration using the steps in [Procedure 17-2](#). It shows how to:

- Create the user **Extremenetworks\_user**, specifying authentication, encryption, and security credentials.
- Assign **Extremenetworks\_user** to the **Extremenetworks** group and associate it to the SNMPv3 security model, **usm**.
- Specify that, if SNMP messages are received with authentication and encryption, the view, **readView** for read requests, and the view **writeView** for write requests will be applied to this user group based on the USM security model.
- For the view **writeView**, include the MIB subtree denoted with OID **1**, and exclude the subtree denoted by OID **1.3.6.1.4.1.5624.1.2.16**.
- Assign an SNMPv3 target parameters entry named **matrixn** to the **Extremenetworks\_user** using the USM security model.
- Create a target address entry named **Extreme\_Networks** at IP address **172.29.10.1** which will use security and authorization criteria contained in a target parameters entry called **matrixn**, and bind these parameters together with a tag entry called **v3TrapTag**.

```
7100 Chassis(su)->set snmp user Extremenetworks_user authentication md5
my_authentication          privacy my_privacy
7100 Chassis(su)->set snmp group Extremenetworks user Extremenetworks_user
security-model usm
7100 Chassis(su)->set snmp access Extremenetworks security-model usm privacy read
readView                   write writeView
7100 Chassis(su)->set snmp view viewname readView subtree 1
7100 Chassis(su)-> set snmp view viewname writeView subtree 1
7100 Chassis(su)-> set snmp view viewname writeView subtree
1.3.6.1.4.1.5624.1.2.16      excluded
7100 Chassis(su)-> set snmp targetparams matrixn user Extremenetworks_user
security-model usm          message-processing v3
7100 Chassis(su)-> set snmp targetaddr Extreme_Networks 172.29.10.1 param matrixn
taglist v3TrapTag
7100 Chassis(su)->set snmp notify SNMPv3TrapGen tag v3TrapTag inform
```

### How SNMP Will Process This Configuration

As described in “[How SNMP Processes a Notification Configuration](#)” on page 17-8, if the SNMP agent on the device needs to send an inform message, it looks to see if there is a notification entry that says what to do with inform messages. Then, it looks to see if the tag list (**v3TrapTag**) specified in the notification entry exists. If it exists, then the inform message is sent to the target addresses specified by the tag list, (**Extreme\_Networks**) using the parameters specified for each address (**matrixn**).

### Configuring an SNMPv3 Inform or Trap Engine ID

This section provides additional information for configuring SNMPv3 inform or trap notifications. The steps in [Procedure 17-3](#) on page 17-14 add to the following configuration example:

```
7100 Chassis(su)->set snmp view viewname All subtree 1
7100 Chassis(su)->set snmp user v3user authentication md5 md5passwd privacy
despasswd
7100 Chassis(su)->set snmp group v3group user v3user security-model usm
7100 Chassis(su)->set snmp access v3group security-model usm privacy exact read
All write All notify All
```

```
7100 Chassis(su)->set snmp notify v3notify tag v3tag inform
7100 Chassis(su)->set snmp targetaddr v3TA 134.141.209.73 param v3TP taglist v3tag
7100 Chassis(su)->set snmp targetparams v3TP user v3user security-model usm
message-processing v3 privacy
```

### Inform EngineIDs

In the Extreme Networks SNMP implementation, the receiver's EngineID value is used by both the sender and receiver to propagate inform notifications. In order to send and receive SNMP v3 informs in their most secure form (with authentication and privacy enabled), you must configure a user ID and corresponding receiver EngineID on the sender as shown in the example in [Procedure 17-3](#). This example assumes that NetSight Console is the receiver, and an 7100-Series switch is the sender.



**Note:** The following file location and EngineID are provided as examples. Your settings will vary.

[Procedure 17-3](#) adds to the configuration example shown in “[Configuring an SNMPv3 Inform or Trap Engine ID](#)” on page 17-13.

### Procedure 17-3 Configuring an EngineID

Step	Task	Command(s)
1.	If necessary, create an SNMP3 configuration.	Refer to “ <a href="#">Configuring an SNMPv3 Inform or Trap Engine ID</a> ” on page 17-13.
2.	On the management station, navigate to and display the Netsight Console SNMP trap configuration file.	<b>C:\Program Files\Extreme Networks\NetSight Shared\snmptrapd.conf</b>
3.	Determine the EngineID from this line in the configuration file.	<b>oldEngineID 0x800007e5804f19000d232aa40</b>
4.	On the Matrix N, define the same user as in the above example ( <b>v3user</b> ) with this EngineID and with the same Auth/Priv passwords you used previously.	<b>set snmp user v3user remote 800007e5804f19000d232aa40 authentication md5 md5passwd privacy despasswd</b>  <div style="display: flex; align-items: flex-start;"> <div style="margin-right: 10px;"></div> <div> <p><b>Note:</b> You can omit the <b>0x</b> from the EngineID. You can also use the colon notation like this: 80:00:07:e5:80:4f:19:00:00:d2:32:aa:40</p> </div> </div>
5.	Navigate to and display the user configuration on the management station. (This assumes that you have already created the user in Netsight Console, so you will only need to add it to the configuration file of the trap daemon.)	<b>C:\Program Files\Extreme Networks\NetSight Console\Bin\snmptrapd.conf</b>
6.	Using any plain text editor, add this line to the configuration file.	<b>createuser v3user MD5 md5passwd DES despasswd</b>

### Trap EngineID

To use traps instead of inform notifications, you would change the preceding configuration as follows:

- Use this command to specify trap notifications:
 

```
set snmp notify v3notify tag v3tag trap
```

- Verify that the “createuser” entry in the NetSight Console SNMP trap configuration looks like this:

```
createuser -e 0x800015f80300e06314d79c v3user MD5 md5passwd DES
despasswd
```

When you are finished modifying the configuration, save the file and restart the SNMP Trap Service using Netsight Services Manager.



**Note:** When installed on a Unix platform, the NetSight server must be manually restarted.

## Configuring an SNMP View

It is possible to include certain OIDs and exclude certain other OIDs within one SNMP MIB view. You do this by stacking different set snmp view includes and excludes which specify a single view name. This allows the user to view all of the “included” OID strings for their associated view name, minus all of the “excluded” OID strings for their view name. If no such parameter is specified, “included” is assumed.

Though it is possible to create and use multiple view names as desired, for demonstration purposes it is simplest to modify the default view, since it is already being referenced by the remainder of the SNMP command set.

The following example removes the default view specifications, and inserts one which permits access to branch MIB **1.3.6.1.2.1** with the exception of branch interfaces **1.3.6.1.2.1.2**:

```
7100 Chassis(su)->clear snmp view All 1
7100 Chassis(su)->clear snmp view All 0.0
7100 Chassis(su)->set snmp view viewname All subtree 1.3.6.1.2.1
7100 Chassis(su)->set snmp view viewname All subtree 1.3.6.1.2.1.2 excluded
7100 Chassis(su)->show snmp view
View Name           = All
Subtree OID         = 1.3.6.1.2.1
Subtree mask        =
View Type           = included
Storage type        = nonVolatile
Row status          = active

View Name           = All
Subtree OID         = 1.3.6.1.2.1.2
Subtree mask        =
View Type           = excluded
Storage type        = nonVolatile
Row status          = active
```

You can test this configuration using any MIB browser directed to the IP of the configured device and using the default community name **public** associated with the view **All**. If configured correctly, only your specified sections of the MIBs will be visible.

## Configuring the Optional Mask Parameter



**Note:** The mechanics of determining exactly how to configure the optional mask parameter make for an inefficient use of time if you will only be using the query once. However, for data retrieved repeatedly, using the method described in the following examples can prevent the unnecessary transfer of much SNMP data over your network.

As defined in RFC2575, an SNMP mask is an optional parameter of the set snmp view command. You can use a mask to modify a view inclusion, designating certain octets of an OID string as wild-card “don't care” values. Once defined, you can view within a MIB branch (using a MIB

browser such as that offered within the NetSight suite of products) only those leaves associated with specific items, such as designated port numbers, MAC addresses, and IP addresses.

For example, the RMON Statistics MIB branch is defined as follows, with the leaves defined within that branch each having multiple iterations, one for each port.

```

etherStatsEntry=1.3.6.1.2.1.16.1.1.1
etherStatsIndex=1.3.6.1.2.1.16.1.1.1.1.<port>
etherStatsDataSource=1.3.6.1.2.1.16.1.1.1.2.<port>
etherStatsDropEvents=1.3.6.1.2.1.16.1.1.1.3.<port>
etherStatsOctets=1.3.6.1.2.1.16.1.1.1.4.<port>
etherStatsPkts=1.3.6.1.2.1.16.1.1.1.5.<port>
etherStatsBroadcastPkts=1.3.6.1.2.1.16.1.1.1.6.<port>
etherStatsMulticastPkts=1.3.6.1.2.1.16.1.1.1.7.<port>
etherStatsCRCAlignErrors=1.3.6.1.2.1.16.1.1.1.8.<port>
etherStatsUndersizePkts=1.3.6.1.2.1.16.1.1.1.9.<port>
etherStatsOversizePkts=1.3.6.1.2.1.16.1.1.1.10.<port>
etherStatsFragments=1.3.6.1.2.1.16.1.1.1.11.<port>
etherStatsJabbers=1.3.6.1.2.1.16.1.1.1.12.<port>
etherStatsCollisions=1.3.6.1.2.1.16.1.1.1.13.<port>
etherStatsPkts64Octets=1.3.6.1.2.1.16.1.1.1.14.<port>
etherStatsPkts65to127Octets=1.3.6.1.2.1.16.1.1.1.15.<port>
etherStatsPkts128to255Octets=1.3.6.1.2.1.16.1.1.1.16.<port>
etherStatsPkts256to511Octets=1.3.6.1.2.1.16.1.1.1.17.<port>
etherStatsPkts512to1023Octets=1.3.6.1.2.1.16.1.1.1.18.<port>
etherStatsPkts1024to1518Octets=1.3.6.1.2.1.16.1.1.1.19.<port>
etherStatsOwner=1.3.6.1.2.1.16.1.1.1.20.<port>
etherStatsStatus=1.3.6.1.2.1.16.1.1.1.21.<port>

```

As shown in the example output above, when displaying the etherStatsEntry branch, all ports are listed for each leaf before moving on to the ports of the next leaf as the result of listing all of the data in numeric OID order.

Here is an abbreviated example of one such SNMP query.

Object	Instance	Type	Value
etherStatsIndex	1001	INTEGER	1001
etherStatsIndex	1518	INTEGER	1518
etherStatsDataSource	1001	OBJECT ID	1.3.6.1...11001
etherStatsDataSource	1518	OBJECT ID	1.3.6.1...12006
etherStatsStatus	1001	INTEGER	valid(1)
etherStatsStatus	1518	INTEGER	valid(1)

## Example

This example shows you how to use the mask parameter to significantly refine your query output, so that only data for specified ports is returned. For this example, assume that 7100-Series slot 1 port 12 is of interest.

The first ten octets of the etherStatsEntry (1.3.6.1.2.1.16.1.1.1) must match exactly as specified. The next octet, representing each of the 21 possible leaves within that branch, need not match exactly. The remainder, representing the port number, must match exactly as specified.

The bit representations for this would be 11111111-11011111, or 0xffdf. If the actual OID string being masked is longer than the specified bits, the missing bits to the right are assumed to be 1's. It is thus only necessary to make the mask long enough (in increments of 8-bit bytes) to designate, with a 0 bit, any desired "wild-card" OID string octets.

The following is an SNMP View using these specifications, starting with a default configuration.

```

7100 Chassis(su)->show snmp view
View Name      = All
Subtree OID    = 1

```

```

Subtree mask      =
View Type         = included
Storage type      = nonVolatile
Row status        = active

```

```

View Name         = All
Subtree OID       = 0.0
Subtree mask      =
View Type         = included
Storage type      = nonVolatile
Row status        = active

```

```

7100 Chassis(su)->clear snmp view All 1
7100 Chassis(su)->set snmp view viewname All subtree 1.3.6.1.2.1.16.1.1.1.0.1012
mask ff:df

```

```

7100 Chassis(su)->show snmp view
View Name         = All
Subtree OID       = 0.0
Subtree mask      =
View Type         = included
Storage type      = nonVolatile
Row status        = active

```

```

View Name         = All
Subtree OID       = 1.3.6.1.2.1.1.1.0.244
Subtree mask      = ff:df
View Type         = included
Storage type      = nonVolatile
Row status        = active

```

You can see by the unexpected Subtree OID value that this view actually accommodates only the right-most 8 bits of the entered decimal value 1012. The hexadecimal equivalent is 0x3f4, and the decimal equivalent of 0xf4 is 244. It is therefore true that this defined subtree will get a “hit” on multiple port values (244, 500, 756, 1012, etc), should they exist. This has nothing to do with the mask, and everything to do with the reasonable limitations of MIB design.



**Note:** Any use of the **mask** parameter assumes the View Type is configured as **included**. Parameters **included** or **excluded** cannot be specified along with the **mask** parameter.

An SNMP query of the etherStatsEntry branch using the community name associated with this defined view would display a result similar to the following.

Object	Instance	Type	Value
etherStatsIndex	1012	INTEGER	1012
etherStatsDataSource	1012	OBJECT ID	1.3.6.1...11012
etherStatsDropEvents	1012	Counter	54323
etherStatsOctets	1012	Counter	302877211
etherStatsPkts	1012	Counter	1592774
etherStatsBroadcastPkts	1012	Counter	793487
etherStatsMulticastPkts	1012	Counter	729406
etherStatsCRCAlignErrors	1012	Counter	0
etherStatsUndersizePkts	1012	Counter	0
etherStatsOversizePkts	1012	Counter	0
etherStatsFragments	1012	Counter	0
etherStatsJabbers	1012	Counter	0
etherStatsCollisions	1012	Counter	0
etherStatsPkts64Octets	1012	Counter	0
etherStatsPkts65to127Octets	1012	Counter	458931
etherStatsPkts128to255Octets	1012	Counter	55190
etherStatsPkts256to511Octets	1012	Counter	656909



etherStatsPkts512to1023Octets	1012	Counter	57
etherStatsPkts1024to1518Octets	1012	Counter	1
etherStatsOwner	1012	OCTET STRING	monitor
etherStatsStatus	1012	INTEGER	valid(1)

## Configuring Secure SNMP Community Names

[Procedure 17-4](#) provides an example of a recommended configuration that will prevent unsecured SNMPv1/v2c access of potentially security compromising information.

As discussed previously in this document, SNMP v1 and v2c are inherently insecure device management protocols. Community names used to define access levels are passed in clear text in all protocol frames sent to the managed entity and may be visible by read-only SNMP users when querying certain SNMP configuration-related objects. In addition, you may be further exposing your network due to configuration conventions which reuse the community names in other aspects of entity management, such as CLI login passwords, and SNMP security names.

Extreme Networks recommends that you “secure” all SNMP community names. You do this by creating a configuration that hides, through the use of “views” sensitive information from SNMP v1/v2c users as follows:

### Procedure 17-4 Configuring Secure Community Names

Step	Task	Command(s)
1.	Create the following SNMP view group configurations. <ul style="list-style-type: none"> <li>An admin (v3) view group with secure read, write, and notify access</li> <li>A read-only view group with unsecure (v1 and v2c) access</li> <li>A read-write view group with unsecure (v1 and v2c) access</li> </ul>	<pre> <b>set snmp access</b> <i>admin-groupname</i> <b>security-model</b> <b>usm</b> <b>privacy exact read</b> <i>secured-viewname</i> <b>write</b> <i>secure-viewname</i> <b>notify</b> <i>secured-viewname</i>  <b>set snmp access</b> <i>read-only-groupname</i> <b>security-model</b> <b>v1 exact read</b> <i>unsecured-viewname</i>  <b>set snmp access</b> <i>read-only-groupname</i> <b>security-model</b> <b>v2c exact read</b> <i>unsecured-viewname</i>  <b>set snmp access</b> <i>read-write-groupname</i> <b>security-model</b> <b>v1 exact read</b> <i>unsecure-viewname</i> <b>write</b> <i>unsecured-viewname</i>  <b>set snmp access</b> <i>read-write-groupname</i> <b>security-model</b> <b>v2c exact read</b> <i>unsecured-viewname</i> <b>write</b> <i>unsecured-viewname</i>           </pre>
2.	Create v1/v2c “public” and “private” community names and security names.	<pre> <b>set snmp community</b> <i>private-communityname</i> <b>securityname</b> <i>read-write-securityname</i>  <b>set snmp community</b> <i>public-communityname</i> <b>securityname</b> <i>read-only-securityname</i>           </pre>

**Procedure 17-4 Configuring Secure Community Names (continued)**

Step	Task	Command(s)
3.	Create user groups and bind them to the security names created in Step 2.	<pre>set snmp group admin-groupname user admin-username  set snmp group read-only-groupname user read-only-securityname security-model v1  set snmp group read-write-groupname user read-write-securityname security-model v1  set snmp group read-only-groupname user read-only-securityname security-model v2c  set snmp group read-write-groupname user read-write-securityname security-model v2c</pre>
4.	Using the <i>admin-username</i> assigned in Step 3, create the v3 user and define authentication keys.	<pre>set snmp user admin-username authentication sha auth-key privacy priv-key</pre>
5.	Using the viewnames assigned in Step 1, create restricted views for v1/v2c users, and unrestricted views for v3 users.	<pre>set snmp view viewname secured-viewname subtree 1  set snmp view viewname secured-viewname subtree 0.0  set snmp view viewname unsecured-viewname subtree 1  set snmp view viewname unsecured-viewname subtree 0.0</pre>
6.	Exclude the following from the restricted view <ul style="list-style-type: none"> <li>• snmpUsmMIB (which contains v3 user names, but no passwords)</li> <li>• snmpVacmMIB (which contains SNMP view configurations)</li> <li>• snmpCommunityTable (which contains community names)</li> </ul>	<pre>set snmp view viewname unsecured-viewname subtree 1.3.6.1.6.3.15 excluded  set snmp view viewname unsecured-viewname subtree 1.3.6.1.6.3.16 excluded  set snmp view viewname unsecured-viewname subtree 1.3.6.1.6.3.18.1.1 excluded</pre>

**Example**

The following example shows an 7100-Series device configuration using the steps in [Procedure 17-4](#).

```
7100 Chassis(su)->set snmp access gAdmin security-model usm privacy exact read
vSecured write vSecured notify vSecured

7100 Chassis(su)->set snmp access gReadOnlyV1V2C security-model v1 exact read
vUnsecured

7100 Chassis(su)->set snmp access gReadOnlyV1V2C security-model v2c exact read
vUnsecured

7100 Chassis(su)->set snmp access gReadWriteV1V2C security-model v1 exact read
vUnsecured write vUnsecured

7100 Chassis(su)->set snmp access gReadWriteV1V2C security-model v2c exact read
vUnsecured write vUnsecured

7100 Chassis(su)->set snmp community cnPrivate securityname sn_v1v2c_rw
7100 Chassis(su)->set snmp community cnPublic securityname sn_v1v2c_ro
7100 Chassis(su)->set snmp group gReadOnlyV1V2C user sn_v1v2c_ro security-model v1
```

```

7100 Chassis(su)->set snmp group gReadWriteV1V2C user sn_v1v2c_rw security-model
v1
7100 Chassis(su)->set snmp group gReadOnlyV1V2C user sn_v1v2c_ro security-model
v2c
7100 Chassis(su)->set snmp group gReadWriteV1V2C user sn_v1v2c_rw security-model
v2c
7100 Chassis(su)->set snmp group gAdmin user it-admin security-model usm
7100 Chassis(su)->set snmp user it-admin authentication sha auth_key privacy
priv_key
7100 Chassis(su)->set snmp view viewname vSecured subtree 1
7100 Chassis(su)->set snmp view viewname vSecured subtree 0.0
7100 Chassis(su)->set snmp view viewname vUnsecured subtree 1
7100 Chassis(su)->set snmp view viewname vUnsecured subtree 0.0
7100 Chassis(su)->set snmp view viewname vUnsecured subtree 1.3.6.1.6.3.15
excluded
7100 Chassis(su)->set snmp view viewname vUnsecured subtree 1.3.6.1.6.3.16
excluded
7100 Chassis(su)->set snmp view viewname vUnsecured subtree 1.3.6.1.6.3.18.1.1
excluded

```

## Reviewing SNMP Settings

Use the **show** commands described in this section to review SNMP settings.

For information about...	Refer to page...
<a href="#">Community</a>	<a href="#">17-20</a>
<a href="#">Context</a>	<a href="#">17-21</a>
<a href="#">Counters</a>	<a href="#">17-21</a>
<a href="#">Engineid</a>	<a href="#">17-22</a>
<a href="#">Groups</a>	<a href="#">17-22</a>
<a href="#">Group Access Rights</a>	<a href="#">17-23</a>
<a href="#">Target Parameter Profiles</a>	<a href="#">17-24</a>
<a href="#">Target Address Profiles</a>	<a href="#">17-24</a>
<a href="#">Notify</a>	<a href="#">17-24</a>
<a href="#">Notify Filter</a>	<a href="#">17-25</a>
<a href="#">Notify Profile</a>	<a href="#">17-25</a>
<a href="#">Users</a>	<a href="#">17-26</a>
<a href="#">Views</a>	<a href="#">17-26</a>

### Community

Use this command to display SNMPv1/SNMPv2c community names and status. In SNMPv1 and v2, community names act as passwords to remote management.

```
show snmp community [name]
```

## Example

```
7100 Chassis(su)->show snmp community public
Name                = public
Security name       = public
Context             =
Transport tag       =
Storage type        = nonVolatile
Status              = active
```

## Context

Use this command to display the context list configuration for SNMP view-based access control:

```
show snmp context
```

## Example

```
7100 Chassis(su)->show snmp context
--- Configured contexts:
default context (all MIBs)
router
```

## Counters

Use this command to display SNMP traffic counter values:

```
show snmp counters
```

## Example

```
7100 Chassis(su)->show snmp counters

--- mib2 SNMP group counters:
snmpInPkts           = 396601
snmpOutPkts          = 396601
snmpInBadVersions    = 0
snmpInBadCommunityNames = 0
snmpInBadCommunityUses = 0
snmpInASNParseErrs   = 0
snmpInTooBigs        = 0
snmpInNoSuchNames    = 0
snmpInBadValues      = 0
snmpInReadOnlys      = 0
snmpInGenErrs        = 0
snmpInTotalReqVars   = 403661
snmpInTotalSetVars   = 534
snmpInGetRequests    = 290
snmpInGetNexts       = 396279
snmpInSetRequests    = 32
```

```

snmpInGetResponses      = 0
snmpInTraps             = 0
snmpOutTooBig          = 0
snmpOutNoSuchNames     = 11
snmpOutBadValues       = 0
snmpOutGenErrs         = 0
snmpOutGetRequests     = 0
snmpOutGetNexts       = 0
snmpOutSetRequests     = 0
snmpOutGetResponses    = 396601
snmpOutTraps           = 0
snmpSilentDrops        = 0
snmpProxyDrops         = 0
--- USM Stats counters:
usmStatsUnsupportedSecLevels = 0
usmStatsNotInTimeWindows    = 0
usmStatsUnknownUserNames    = 0
usmStatsUnknownEngineIDs    = 0
usmStatsWrongDigests        = 0
usmStatsDecryptionErrors     = 0

```

## Engineid

Use this command to display SNMP engine properties:

```
show snmp engineid
```

### Example

```

7100 Chassis(su)->show snmp engineid

EngineId: 80:00:15:f8:03:00:e0:63:9d:b5:87
Engine Boots      = 12
Engine Time       = 162181
Max Msg Size      = 2048

```

## Groups

Use this command to display SNMP group information. If no parameters are specified, all information about all groups is displayed.

```
show snmp group [groupname groupname] [user user] [security-model {v1 | v2c | usm}] [volatile | nonvolatile | read-only]
```

### Example

```

7100 Chassis(su)->show snmp group
Security model      = SNMPv1
Security/user name = public
Group name         = groupRW

```

```

Storage type           = nonVolatile
Row status             = active

Security model         = SNMPv2c
Security/user name    = public
Group name             = groupRW
Storage type          = nonVolatile
Row status            = active

Security model         = USM
Security/user name    = admin1
Group name            = alladmin
Storage type          = nonVolatile
Row status            = active

Security model         = USM
Security/user name    = admin2
Group name            = alladmin
Storage type          = nonVolatile
Row status            = active

```

## Group Access Rights

Use this command to display an SNMP group's access rights. If no parameters are entered, access information about all groups is displayed.

```

show snmp access [groupname] [security-model {v1 | v2c | usm}]
[noauthentication | authentication | privacy] [context context] [volatile |
nonvolatile | read-only]

```

### Example

```

7100 Chassis(su)->show snmp access
Group           = groupRW
Security model  = SNMPv1
Security level  = noAuthNoPriv
Read View       = All
Write View      = All
Notify View     = All
Context match   = "default context" (exact)
Storage type    = nonVolatile
Row status      = active

Group           = groupRW
Security model  = SNMPv2c
Security level  = noAuthNoPriv
Read View       = All
Write View      = All

```

```
Notify View      = All
Context match    = "default context" (exact)
Storage type     = nonVolatile
Row status       = active
```

## Target Parameter Profiles

Use this command to display SNMP target parameter profiles. If no parameters are specified, information for all target parameter profiles is displayed.

```
show snmp targetparams [targetParams] [volatile | nonvolatile | read-only]
```

### Example

```
7100 Chassis(su)-> show snmp targetparams matrix

Target Parameter Name = matrix
Security Name          = Extremenetworks_user
Message Proc. Model   = USM
Security Level         = authNoPriv
Storage type           = nonVolatile
Row status             = active
```

## Target Address Profiles

Use this command to display SNMP target address information. If no parameters are entered, information about all target address profiles is displayed.

```
show snmp targetaddr [targetAddr] [volatile | nonvolatile | read-only]
```

### Example

```
7100 Chassis(su)-> show snmp targetaddr

Target Address Name = Extremenetworks_user
Tag List           =
IP Address          = 172.29.10.1
UDP Port#           = 162
Target Mask         = 255.255.255.255
Timeout             = 1500
Retry count         = 3
Parameters          = matrix
Storage type        = nonVolatile
Row status          = active
```

## Notify

Use this command to display the SNMP notify configuration, which determines which management targets will receive SNMP notifications. If no parameters are entered, information about all notify configurations is displayed.

```
show snmp notify [notify] [volatile | nonvolatile | read-only]
```

## Example

```
7100 Chassis(su)->show snmp notify
Notify name      = 1
Notify Tag       = Console
Notify Type      = trap
Storage type     = nonVolatile
Status           = active

Notify name      = 2
Notify Tag       = TrapSink
Notify Type      = trap
Storage type     = nonVolatile
Status           = active
```

## Notify Filter

Use this command to display SNMP notify filter information, identifying which profiles will not receive SNMP notifications:

```
show snmp notifyfilter [profile] [subtree oid-or-mibobject] [volatile |
nonvolatile | read-only]
```

## Example

```
7100 Chassis(su)->show snmp notifyfilter
Profile          = pilot1
Subtree          = 1.3.6
Subtree mask
Filter type      = included
Storage type     = nonVolatile
Row status       = active
```

## Notify Profile

Use this command to display SNMP notify profile information:

```
show snmp notifyprofile [profile] [targetparam targetparam] [volatile |
nonvolatile | read-only]
```

## Example

```
7100 Chassis(su)->show snmp notifyprofile area51
Notify Profile   = area51
TargetParam      = v3ExampleParams
Storage type     = nonVolatile
Row status       = active
```



## Users

Use this command to display SNMPv3 users:

```
show snmp user [list] | [user] | [remote remote] [volatile | nonvolatile | read-only]
```

### Example

```
7100 Chassis(su)->show snmp user Extremeretworks_user
```

```
EngineId           = 80:00:15:f8:03:00:e0:63:9d:cb:89
Username           = Extremeretworks_user
Auth protocol      = usmHMACMD5AuthProtocol
Privacy protocol   = usmDESPrivProtocol
Storage type       = nonVolatile
Row status         = active
```

## Views

Use this command to display SNMP views. If no parameters are entered, all view information is displayed.

```
show snmp view [viewname] [subtree oid-or-mibobject] [volatile | nonvolatile | read-only]
```

### Example

```
7100 Chassis(su)->show snmp view readView
```

```
View Name          = readView
Subtree OID        = 1
Subtree mask       =
View Type          = included
Storage type       = nonVolatile
Row status         = active
```

## Spanning Tree Configuration

This chapter provides the following information about configuring and monitoring Spanning Tree protocols on Extreme Networks 7100-Series devices:

For information about...	Refer to page...
<a href="#">What Is the Spanning Tree Protocol?</a>	<a href="#">18-1</a>
<a href="#">Why Would I Use Spanning Trees in My Network?</a>	<a href="#">18-2</a>
<a href="#">How Do I Implement Spanning Trees?</a>	<a href="#">18-3</a>
<a href="#">STP Overview</a>	<a href="#">18-3</a>
<a href="#">Functions and Features Supported on the 7100-Series Device</a>	<a href="#">18-6</a>
<a href="#">Understanding How Spanning Tree Operates</a>	<a href="#">18-9</a>
<a href="#">Configuring STP and RSTP</a>	<a href="#">18-21</a>
<a href="#">Configuring MSTP</a>	<a href="#">18-27</a>
<a href="#">Understanding and Configuring SpanGuard</a>	<a href="#">18-32</a>
<a href="#">Understanding and Configuring Loop Protect</a>	<a href="#">18-34</a>
<a href="#">Terms and Definitions</a>	<a href="#">18-39</a>

### What Is the Spanning Tree Protocol?

The Spanning Tree Protocol (STP) resolves the problem of physical loops in a network by establishing one primary path between any two devices. Duplicate paths are barred from use and become standby or “blocked” paths until the primary path fails, at which point the redundant path can be brought into service.

STP operates by forming a fully connected tree of data loop free LAN connected bridges (switches) through the exchange of Bridge Protocol Data Units (BPDUs). Each bridge port transmits BPDUs on a periodic basis. The information contained in the BPDU is used by the receiving bridge to calculate a port role for each bridge port. There is one bridge in the network chosen to be the root bridge, based on its bridge ID. Ports that directly connect bridges to the root bridge or are connected through another bridge are assigned one of four roles:

- Root Port – The best path to the root
- Designated Port – Ports which either provide a path to the root for other bridges or connect end users
- Backup Port – A port attached to a LAN where another port of the same bridge is a designated port. This backup port takes over the designated role should the LAN’s designated port become disabled

- Alternate Port – Port providing a path to the root that is not root, designated, or backup

For a summary of port roles, see [Table 18-1](#) on page 18-14.

While the network is in a steady state, alternate and backup ports are in blocking state; root and designated ports are in forwarding state.


STP allows for the automatic reconfiguration of the network. When bridges are added to or removed from the network, root election takes place and port roles are recalculated.

## Why Would I Use Spanning Trees in My Network?

Redundant links must be factored into even the simplest of topologies to protect against data loss and downtime due to any single point of failure. STP prevents redundant links from forming data loops which would consume all available network bandwidth. STP manages redundant links by keeping them in a blocking state and automatically unblocking them when changes in topology require that they be used. See [Table 18-2](#) on page 18-14 for a summary of Spanning Tree port states.


As shown in [Figure 18-1](#), a planned redundant link between Switch 1 and Switch 2 makes it possible for a bridging loop to occur. If Station 1 transmits a multicast or broadcast packet to Station 2 in this scenario, the packet would continue to circulate endlessly between both switching devices. Without Spanning Tree blocking one of the links, there would be nothing at layer 2 to stop this loop from happening and unnecessarily consuming network resources. As administrator, you would be forced to manually disable one of the links between Switch 1 and 2 for the [Figure 18-1](#) network to operate.

**Figure 18-1 Redundant Link Causes a Loop in a Non-STP Network**



STP automatically blocks redundant paths, as shown in [Figure 18-2](#). In the event that the primary (unblocked) path fails, STP places the blocked path into service. If the original primary path recovers, the redundant path will once again block and the primary path will be used.

**Figure 18-2 Loop Avoided When STP Blocks a Duplicate Path**



# How Do I Implement Spanning Trees?

By default, Spanning Tree is both enabled globally and on all ports. The design of the Spanning Tree protocol and the default configuration values on these devices make user configuration unnecessary in order to add redundant ports to your network. You will want to make configuration changes to select a root bridge, take advantage of Multiple Spanning Tree, or use any of the advanced features described below. Before configuring STP it is important to understand how it works.

## STP Overview

Extreme Networks switch devices support the Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP) as defined in the following standards and described in IEEE 802.1Q:

- IEEE 802.1D (Spanning Tree Protocol)
- IEEE 802.1w (Rapid Spanning Tree Protocol)
- IEEE 802.1s (Multiple Spanning Tree Protocol)
- IEEE 802.1t (Update to 802.1D)

STP forms a network of bridges connected by LANs into a tree that is:

- Predictable – A given set of configured bridges always yields the same topology when the network reaches steady state
- Optimized – STP selects the best path to the root bridge
- Fully connected – Each bridge communicates with every other bridge in the network
- Free from data loops – One root port is chosen in a bridge and the remaining ports with paths to the root bridge are put into blocking mode

The root bridge is the bridge with the lowest bridge ID in the network and functions as the logical center of the STP network. Each bridge calculates its best path to the root using the information contained in BPDUs received from its neighbor bridges. Non-root bridges select the root port among all the ports receiving BPDUs. BPDUs advertise a bridge's cost to the root bridge. The root port is chosen from the ports with received BPDUs indicating a path to the root. The root port will have the lowest cost path to the root. In the case of multiple ports offering identical costs, tie breaking is based upon the transmitting bridge ID, transmitting port ID, and receive port ID. For MSTP there are additional fields to consider – internal path cost and regional root ID. These are all discussed in more detail below.

Once the root port has been established, STP determines the other port roles. Ports providing a path to root but are not the root port become alternate ports because they provide an alternate path to the root. Other operational ports that provide a path to the root for attached bridges have the designated role. There is another type of port known as a backup port. A backup port attaches to a LAN where another port of the same bridge is a designated port. A backup port does not become part of the active topology unless the LAN's designated port is disabled and the backup port takes over the designated role.

The alternate and backup ports are set to blocking state while the root and designated ports move to the forwarding state.

Bridge priority, port path cost, and port priority are configurable parameters that are part of the port role calculation and may be modified to create the desired topology.

**Bridge Priority** – A typical network configuration would place two or more bridges in the core. To preserve root in the core, the core bridges would each have their bridge priority set to a lower value than bridges you do not desire to be root. The default bridge priority value is 32768. If you

desire a particular bridge to be root, set its bridge priority to a lower value than bridges that should not be root. Otherwise the bridge with the lowest MAC address is set to root.

**Port Path Cost** – If it is desired for a bridge to use one link over another, the administrative port path cost may be modified. The default of zero ensures that the link with the highest speed gets chosen.

**Port Priority** – Port Priority may be set but is not typically modified, as Link Aggregation is usually run on multiple links between two bridges.

## Rapid Spanning Tree

Rapid Spanning Tree (RSTP) optimizes convergence in a properly configured network by significantly reducing the time to reconfigure the network's active topology when physical topology or configuration parameter changes occur. RSTP is defined in the IEEE 802.1w standard. Spanning Tree's primary goal is to ensure a fully connected, loop-free topology. A secondary goal, realized with the introduction of RSTP, is to move root and designated ports to the forwarding state as quickly as possible.

In a stable topology all the root and designated ports will be forwarding and the alternate and backup ports will be blocking. When there is a network topology change, Spanning Tree recalculates port roles. Ports which are no longer part of the active topology will be put into blocking state. New designated ports will only forward after receiving an acknowledgement or, in the case of a port being connected to a non-RSTP device (802.1d), after a sufficient amount of time has passed.

When a topology change occurs, a change in port operational status or new information contained in BPDUs is immediately acted upon. A new root port moves to forwarding state as soon as any recent former root port is put into blocking state. A designated port moves to forwarding state once the connected device acknowledges agreement with the new topology information. This is typically an exchange of two BPDUs. These rules ensure an orderly transition from the old topology to the new topology by preventing transient loops.

## Multiple Spanning Tree

The Multiple Spanning Tree Protocol (MSTP) provides for traffic forwarding on multiple ports for each bridge. A single Spanning Tree only allows for single root port forwarding per bridge. MSTP provides for a number of common network requirements that cannot be configured on a single Spanning Tree (for example, the segregation of traffic over multiple VLANs or optimizing the utilization of redundant links between switching devices in a network).

An MSTP configuration is made up of one or more:

- Multiple Spanning Tree (MST) Regions – A set of connected bridges that share the same MST configuration ID
- MST configuration IDs – A unique identifier for each MST region
- Spanning Tree Identifiers (STIs) – A unique identifier for each Spanning Tree

An MSTP configuration is made up of zero or more Multiple Spanning Tree Instances (MSTIs). An MSTI is an STI that exists within an MST region other than the default STI 0.

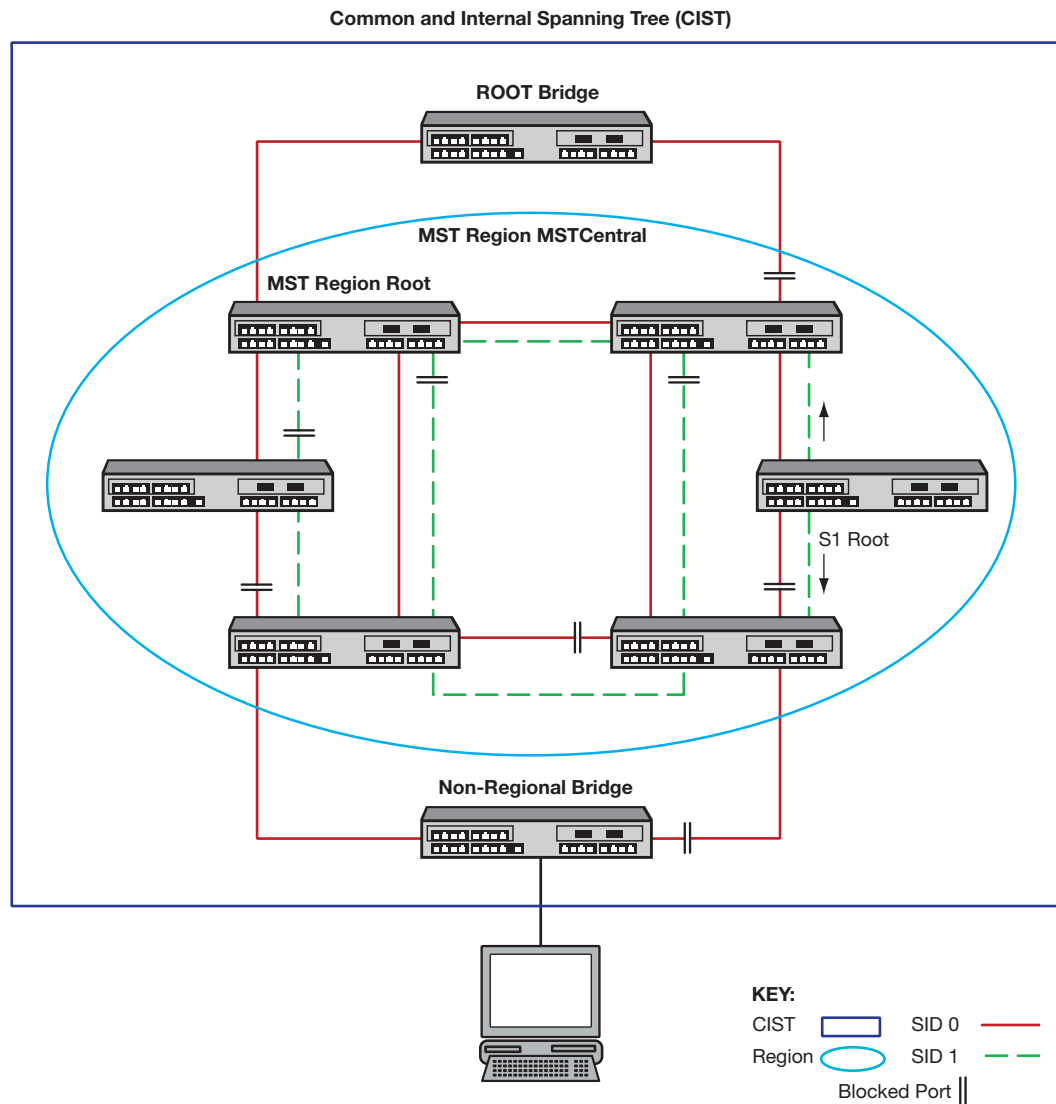
All bridges in the Spanning Tree network are inter-connected by STI 0 and can belong to:

- The Common Spanning Tree (CST) – A Spanning Tree defined in the IEEE 802.1q standard that assumes one Spanning Tree instance for the entire bridged network, regardless of the number of VLANs

- An Internal Spanning Tree (IST) instance – A Spanning Tree instance that extends the CST inside the MST region and represents the entire MST region as a single CST virtual bridge to the outside world.

One or more MSTs can be part of the Common and Internal Spanning Tree (CIST). The CIST represents the connectivity of the entire network. [Figure 18-3](#) provides an overview of an MST configuration with one MST region within the CIST. The MST region's configuration ID name is **MSTCentral**.

**Figure 18-3 Multiple Spanning Tree Overview**



SID 0 is the default Spanning Tree and interconnects all bridges to the Root Bridge. SID 0 within the MST is the Internal Spanning Tree (IST) and provides connectivity out to the CST as well as functioning as another Spanning Tree instance within the MST region. SID 1 is an MSTI configured within the MST region.

Each SID has a root bridge. In [Figure 18-3](#) the SID 0 root bridge belongs to the CST. The SID 0 root bridge functions as root for SID 0 Spanning Tree instance in both the CST and MST. SID 1 only exists within the **MSTCentral** region. The root for SID 1 is a bridge within the MSTCentral region. SID 1 can provide traffic segmentation by forwarding traffic on a second VLAN within the

MSTCentral region or provide for optimization of redundant links by forwarding traffic within the MSTCentral region on the same VLAN.

See [Configuring MSTP](#) on page 18-27 for examples of MSTP traffic segregation and optimization of redundant links.



**Note:** MSTP and RSTP are fully compatible and interoperable with each other and with legacy STP.

## Functions and Features Supported on the 7100-Series Device

### Spanning Tree Versions

MSTP and RSTP automatically detect the version of Spanning Tree being used on a LAN. RSTP bridges receiving MSTP BPDUs interpret them as RSTP BPDUs. MSTP and RSTP bridges receiving STP BPDUs will switch to use STP BPDUs when sending on the port connected to the STP bridge. MSTP incorporates a force version feature that allows you to administratively force MSTP to behave as STP or RSTP. This will cause all ports of the bridge to transmit STP or RSTP BPDUs. Use the force version feature when the MSTP bridge is attached to a device that cannot properly handle a non-STP BPDU.



**Note:** Forcing a bridge to STP will prevent it from joining a region and will disable rapid reconfiguration.

### Maximum SID Capacities

By default, Multiple Spanning Tree mode is globally enabled on Extreme Networks switching devices and a single Spanning Tree is configured as SID 0.

Maximum device SID capacities in addition to SID 0 are instances.

### Network Diameter

Extreme Networks switching devices support a default 20-bridge span from and including the root bridge. You can configure support for a maximum span of up to 40 bridges from the Spanning Tree root in the Common Spanning Tree (CST) or the Common and Internal Spanning Tree (CIST) regional root within an MST region. Max age defines the diameter for the CST and Maxhops defines the diameter within a region. See [Defining the Maximum Age Time](#) on page 18-25.

### Port Forwarding

MSTP and RSTP use rapid forwarding mechanisms to get ports to the forwarding state. However, there is a difference in forwarding time between user ports and inter-switch links (ISLs). If a user port is defined as `admedge TRUE` using the `set spantree admedge` command, it will forward as soon as the port becomes operational. An ISL will forward based on an exchange of BPDUs. By default, `autoedge` is set to `TRUE` and `admedge` is set to `FALSE`. These settings satisfy most requirements. `Autoedge` allows a port defined as `admedge FALSE` to discover in a short period of time that it is an edge port. The only time it is necessary to set `admedge` to `TRUE` is when the attached user device cannot tolerate the several seconds required for auto-detection to detect the

port as a user port and move it to forwarding. Setting an ISL to `admindedge TRUE` should be avoided because it can lead to transient data loops.

## Disabling Spanning Tree

Spanning Tree may be disabled globally or on a per port basis. If Spanning Tree is disabled globally all linked ports will be in a forwarding state and the Spanning Tree Protocol will not run. Additionally, a received BPDU will be treated as any multicast packet and flooded out all ports.

If Spanning Tree is disabled on a port by setting `portadmin` to `disabled` using the `set spantree portadmin` command, the port will be in a forwarding state and the protocol will not run for that port. A received BPDU will be consumed. The intention is that the port terminates the Spanning Tree domain. For instance, the port may be attached to a router. If this port were accidentally attached to another switching port, a data loop may result.

## STP Features

Extreme Networks switching devices provide seamless Spanning Tree functionality by:

- Creating a single Spanning Tree from any arrangement of switching or bridging elements.
- Compensating automatically for the failure, removal, or addition of any switching device in an active data path.
- Achieving port changes in short time intervals, which establishes a stable active topology quickly with minimal network disturbance.
- Using a minimum amount of communications bandwidth to accomplish the operation of the Spanning Tree Protocol.
- Reconfiguring the active topology in a manner that is transparent to stations transmitting and receiving data packets.
- Managing the topology in a consistent and reproducible manner through the use of Spanning Tree Protocol parameters.
- Increasing security and reliability with SpanGuard, as described below and in [Understanding and Configuring SpanGuard](#) on page 18-32.
- Further protecting your network from loop formation with Loop Protect, as described below and in [Understanding and Configuring Loop Protect](#) on page 18-34.
- Supporting more port density and faster port speeds as described in [Updated 802.1t](#) on page 18-8.
- Supporting the Restricted Topology Change Notice (TCN) feature as described in [Restricted Topology Change Notification \(TCN\)](#) on page 18-8.
- Supporting the Restricted Role feature as described in [Restricted Role](#) on page 18-9

## SpanGuard

The Extreme Networks SpanGuard feature helps protect your network from two situations that can cause a Denial of Service (DoS) condition: repeated topology change notifications and an unwanted bridge being inserted into and forcing traffic through the topology. SpanGuard increases security and reliability by preventing Spanning Tree respans that can occur when BPDUs are received on user ports and notifies network management that they were attempted.

If a SpanGuard enabled port receives a BPDU, it becomes locked and transitions to the blocking state. It will only transition out of the blocking state after a globally specified time or when it is manually unlocked. By default, SpanGuard is globally disabled on the 7100-Series device and



must be globally enabled to operate on all user ports. For a more detailed discussion of the SpanGuard feature, refer to [Understanding and Configuring SpanGuard](#) on page 18-32.

## Loop Protect

The Loop Protect feature prevents or short circuits loop formation caused by redundant paths in your network by requiring ports to receive BPDUs (RSTP/MSTP only) on point-to-point ISLs before their states are allowed to become forwarding. Further, if a BPDU timeout occurs on a port, its state becomes listening until a new BPDU is received. In this way, both upstream and downstream facing ports are protected.

When a root or alternate port loses its path to the root bridge, due to message age expiration, it takes on the role of designated port and will not forward traffic until a BPDU is received. When a port is intended to be the designated port in an ISL, it constantly proposes and will not forward until a BPDU is received. It will revert to listening if it stops getting a response. Loop Protect also overrides the port admin setting. This protects against misconfiguration (such as disabling STP on a port using the `set spantree portadmin port-string disable` command) and protocol failure by the connected bridge. By default, the Loop Protect feature is globally disabled on Extreme Networks switch devices and must be globally enabled to operate on all ports. For configuration information, refer to [Understanding and Configuring Loop Protect](#) on page 18-34.

## Updated 802.1t

IEEE 802.1t is enabled by default on Extreme Networks switch devices. This updated Spanning Tree protocol supports multiple Spanning Trees, more switch port density, and faster port speeds.

802.1t includes the following updates:

- New bridge identifier encoding (4-bit priority, 12-bit system ID extension, 48-bit bridge address)
- New port identifier encoding (4-bit priority, 12-bit port number)
- Bridge detection state machine (for edge port identification)
- Path cost default values (the ability to switch between 802.1t and 802.1d mode and cost values)

## Restricted Topology Change Notification (TCN)

Restricted Topology Change Notification (TCN) is a Spanning Tree protocol feature that allows or disallows TCN propagation on specified ports. When Restricted TCN is disabled, TCN propagation is allowed. The port propagates received TCNs and topology changes to other ports. Restricted TCN is disabled by default. When Restricted TCN is enabled, the port does not propagate received TCNs and topology changes to other ports. Enable Restricted TCN to prevent unnecessary address flushing in the core region of the network caused by activation of bridges external to the core network.

A possible reason for not allowing TCN propagation is when bridges are not under the full control of the administrator or because MAC operational state for the attached or downstream LANs transitions frequently, causing disruption throughout the network.

Rapid Spanning Tree responds to TCNs by selectively flushing the filter database. Persistent TCNs are disruptive, causing persistent address flushing, which in turn causes increased flooding in the network. Restricted TCN is a useful tool when it is not possible to remove the source of the TCNs. Be cautioned that when enabled, temporary loss of connectivity can occur after changes in a Spanning Tree's active topology, due to persistent, incorrectly learned, station location information. This would be the case where the part of the topology that is the source of the unwanted TCNs is redundantly connected to other parts of the network.

## Restricted Role

Restricted Role is a Spanning Tree protocol feature that allows or disallows the root role on specified ports. When Restricted Role is enabled, the port will not be selected as the root port for the CIST or any MSTI, even if it has the best Spanning Tree priority. A port with Restricted Role enabled is selected as an alternate port after the root port has been selected.

If enabled, Restricted Role can cause lack of Spanning Tree connectivity. Setting Restricted Role to enabled prevents bridges, external to a core region of the network, from influencing the Spanning Tree active topology. You may wish to use Restricted Role when bridges are not under your full control. You may also wish to enable Restricted Role on ports where the bridge is external to the core and where the port faces away from the root, in cases where the port role would normally be designated. This can speed network reconvergence, particularly after loss of the root bridge. Restricted role is disabled by default.

## Multisource Detection

Multisource detection is a feature that prevents network disruption due to excessive topology changes caused by a full duplex port transmitting multiple BPDUs with different source MAC addresses, and hence different BPDU information.

When a port is point-to-point, the received priority information comes from the most recently received BPDU. When a port is non-point-to-point, the received information reflects the best priority information out of all the received BPDUs. Typical scenarios for multisource detection are when a switch is connected to a device which

- has been improperly configured to forward received BPDUs out other ports, or
- has been configured to not run the Spanning Tree protocol and treats BPDUs as multicast packets by transmitting them out all other forwarding ports.

In these situations, the connected port is effectively acting as a shared media device. The way to detect shared media is the duplex setting. Since the port is full duplex, it treats the connection as point-to-point. Multisource detection, which is always enabled, recognizes the multiple source MAC addresses and sets the port's operational point-to-point status to false, treating the port as a shared media device. The port is constantly monitored. If the situation is resolved, as determined by receiving a unique address for a sufficient amount of time, the port's operational point-to-point status will revert to true.

A syslog message is issued when multiple source addresses are detected.



**Note:** When loop protect is configured for the port, if multisource detection is triggered, the port will go to the listening state and no longer be part of the active topology. Loop protect does not operate on shared media ports.

## Understanding How Spanning Tree Operates

This section provides you with a more detailed understanding of how the Spanning Tree operates in a typical network environment.

For information about...	Refer to page...
<a href="#">Spanning Tree Basics</a>	18-10
<a href="#">Electing the Root Bridge</a>	18-10
<a href="#">Assigning Path Costs</a>	18-10
<a href="#">Paths to Root</a>	18-10

For information about...	Refer to page...
<a href="#">Identifying Designated, Alternate, and Backup Port Roles</a>	18-12
<a href="#">Assigning Port States</a>	18-14
<a href="#">RSTP Operation</a>	18-14
<a href="#">MSTP Operation</a>	18-15
<a href="#">Multisource Detection</a>	18-20

## Spanning Tree Basics

The most elemental task of a Spanning Tree Bridge is to control the forwarding state of each port. The bridge evaluates the information received from its immediate neighbors in the form of BPDUs, along with its own configured information. From this information a root is elected and then port roles may be selected for each port. For the root port and designated ports the desired state is forwarding. These ports will become forwarding by subsequent exchange of BPDUs or through the expiration of protocol timers according to the state machines defined by the Spanning Tree Protocol. The remaining ports will become discarding (shorthand for the states of blocking, listening, and learning).

To facilitate this process, the bridge transmits BPDUs out each port on a periodic basis as well as in response to events such as changes in port operational status, configuration changes, timer expiration, and changes in topology derived from received BPDUs.

## Electing the Root Bridge

The network topology is determined by the selection of the root bridge. The topology is based on each bridge's best path to root. Root election occurs on each bridge when new information is received from a neighboring bridge in a BPDU, when link is lost on a port connecting a neighboring bridge, or when the bridge's priority is administratively changed.

The root is elected by comparing the root IDs received in BPDUs as well as the bridge's own bridge ID. The bridge with the lowest ID is chosen as root. The bridge ID is an 8-byte value with the 2 most significant bytes being the bridge priority and the 6 least significant bytes being the bridge MAC address. Root may be forced to a particular bridge by the configuration of bridge priority. Among bridges with the same bridge priority, the one with the lowest MAC address is elected root. If a bridge receives no BPDUs indicating a better bridge ID than its own, it becomes the root bridge.

## Assigning Path Costs

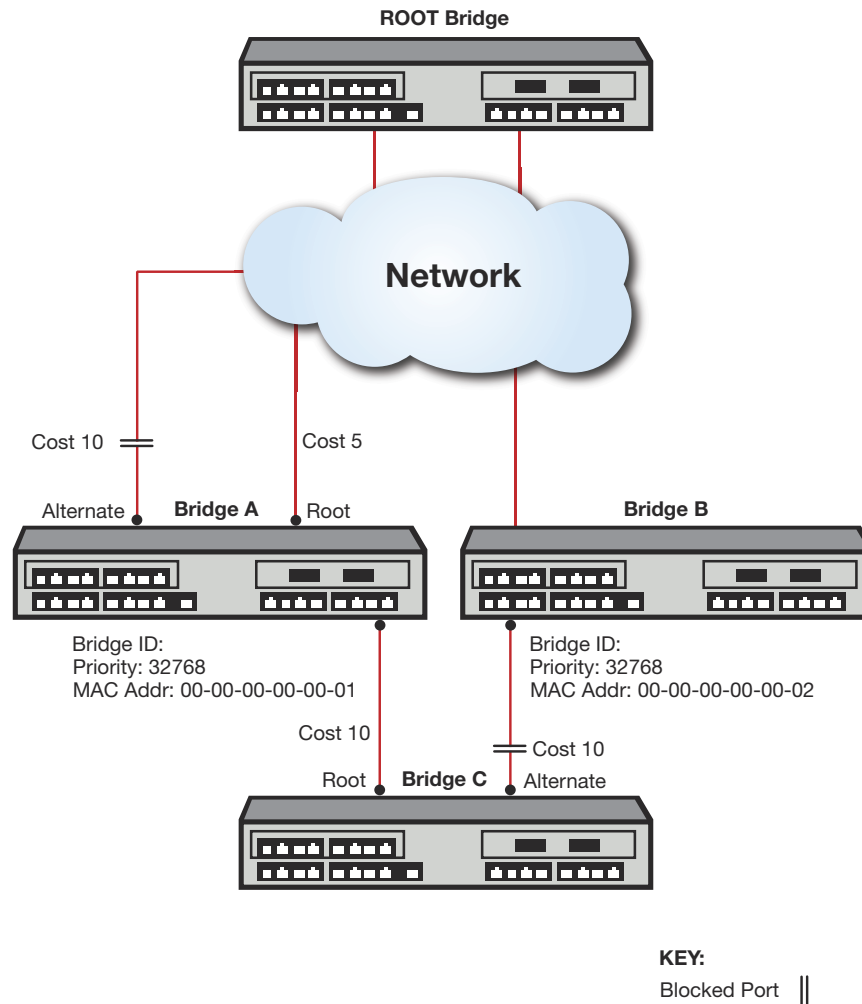
Path costs are one factor in determining port roles. Each LAN segment has an operational path cost associated with it. The cost is based on the port speed, by default. The higher the speed, the lower the cost. Port costs for link aggregations are based on the aggregate speed of all the underlying physical ports. The port cost value may also be administratively assigned using the `set spantree adminpathcost` command. This may be done to choose a particular path.

## Paths to Root

If the bridge is not elected as root, one or more ports provide a path back to the root bridge. The port with the best path is selected as the root port. The best path is the one that has the lowest designated cost. The lowest cost is the aggregate cost of all the LANs traversed between the port and the root bridge. [Figure 18-4](#) on page 18-11 displays root port configuration based upon lowest


cost for Bridge A. If multiple ports have the same lowest cost, the one with the lowest bridge ID becomes the root port. The bridge ID is the ID of the transmitting bridge. [Figure 18-4](#) displays root port configuration based upon lowest bridge ID for Bridge C.

**Figure 18-4 Root Port Selection Based Upon Lowest Cost or Bridge ID**



If there are ports with the same bridge ID, the port ID is used as a tie breaker. The port with the lowest port ID is chosen as root port. The port ID is a 2-byte value with the 4 most significant bits being the port priority and the 12 least significant bits being the bridge port number. Because the port priority occupies the most significant bits in the port ID, setting a lower port priority assures that port will be selected as root. In the case of no single port having a lowest port priority, the root port is selected based upon the overall port ID value. [Figure 18-5](#) on page 12 presents a root port configuration for Bridge B determined by the port priority setting. If there is still a tie, these ports are connected via a shared medium. The final tie breaker is the receiving port ID.


**Figure 18-5 Root Port Selection Base Upon Lowest Port ID**



After selecting the root port, the bridge's cost to root is the total of the root port's designated cost as advertised in the received BPDU, plus the path cost associated with that port. In a hierarchically designed network, the designated cost for ports attached to the next higher level will be less than the bridge's cost to the root. Ports attached to bridges on the same level will have designated costs equal to the bridge's cost to the root. Likewise, this bridge's port will advertise the bridge's cost to the root. Thus one port connected to the LAN will be the designated port and the other(s) will be alternate. The port on the bridge with the lower ID will be the designated port.

### Identifying Designated, Alternate, and Backup Port Roles

Ports in a Spanning Tree configuration are assigned one of four roles: root, designated, alternate, or backup. [Figure 18-6](#) on page 18-13 presents an overview of Spanning Tree port roles.

**Figure 18-6 Spanning Tree Port Role Overview**

- |   |                                   |   |                                   |
|---|-----------------------------------|---|-----------------------------------|
| 1 | Port 1, Bridge A, Alternate Port  | 5 | Port 5, Bridge A, Backup Port     |
| 2 | Port 2, Bridge A, Root Port       | 6 | Port 6, Bridge A, Backup Port     |
| 3 | Port 3, Bridge A, Designated Port | 7 | Port 7, Bridge A, Designated Port |
| 4 | Port 4, Bridge A, Designated Port | 8 | Port 1, Bridge B, Root Port       |

All ports which act as edge ports take on the designated port role. If the bridge has been elected root, all ports connected to ports on other bridges are also designated ports.

On non-root bridges, Spanning Tree identifies ports which provide a path to the root bridge and selects the best path among these as the root port as described in [Paths to Root](#) on page 18-10 ([Figure 18-6](#), callout2). There may be only a single port providing a path to root, in which case that is the root port and the remaining ports are designated. If there are other ports providing a path to root, these ports are selected as alternate paths. Should the root port become disabled, one of the alternate ports will be selected as the new root port ([Figure 18-6](#), callout 1).

A port which is not a designated port, but is connected to another port on the same bridge ([Figure 18-6](#), callout 5) or connected to a shared LAN on which this bridge already provides a designated port ([Figure 18-6](#), callout 6), takes the role of backup port. In the shared LAN example it may take over as designated port if the original designated port is disabled.

All operational ports which are not root, alternate or backup are designated ports. These ports provide a path to the root for attached devices.

[Table 18-1](#) on page 18-14 provides a summary of STP port roles.

**Table 18-1 Spanning Tree Port Roles**

Port Role	Description
Root	The one port that is used to connect to the root bridge. It is elected based on its least “path-cost” to the root bridge and is forwarding traffic.
Alternate	Any redundant upstream port that provides an alternate path to the root bridge (other than the root port). Alternate ports are set to blocking.
Designated	Any downstream port that provides a path back to the root bridge for a downstream bridge. This port is forwarding traffic.
Backup	A port that acts as a redundant designated port on a shared LAN. Backup ports are set to blocking.

## Assigning Port States

All ports are blocking when the operational status switches from disabled to enabled. By default, automatic edge detection is enabled and ports are configured as non-edge ports. In this scenario a user port will become forwarding in several seconds. A port configured as an edge port will forward immediately.

Ports which are selected as alternate or backup ports are immediately put into the discarding state and remain discarding until a new port role is selected. The root port may go to the forwarding state as long as any recent former root ports are synchronized with the new root information. Designated ports may forward as soon as the attached port signals agreement as specified by RSTP. In the absence of the above conditions, root and designated ports get to the forwarding state through the use of timers. The value of the timers is dependent on the value of ForceVersion. The default value is MSTP. If the value is StpCompatible, the timer values are derived from forward delay. Otherwise the values are derived from hello time.

[Table 18-2](#) provides a summary of STP port states.

**Table 18-2 Spanning Tree Port States**

Port State	Behavior
Blocking	Actively preventing traffic from using this path. Still receiving BPDUs, so continuing to monitor for management and STA information.
Listening	Continuing to block traffic while waiting for protocol information to determine whether to go back to the blocking state or continue to the learning state. Listens to BPDUs to ensure no loops occur on the network.
Learning	Learning station location information but continuing to block traffic.
Forwarding	Forwarding traffic and continuing to learn station location information.
Disabled	Disabled administratively or by failure.
Discarding	Used as shorthand for blocking, listening, or learning state.

## RSTP Operation

RSTP optimizes convergence by significantly reducing the time to reconfigure the network’s active topology when physical topology or configuration parameter changes occur. RSTP provides rapid connectivity following the failure of a switching device, switch port, or the addition of a switch into the network.

A new root port may forward as soon as any recent root ports are put into blocking.

A designated port may forward with the exchange of two BPDUs in rapid succession. The designated port presents new BPDU information with a proposal request. The attached port processes the BPDU and may respond immediately with an agreement. Upon reception of that agreement BPDU, the designated port may move to forwarding. Another feature of RSTP is that designated ports transmit periodic BPDUs regardless of reception of BPDUs at the root port. This insulates the network from jitter in receiving BPDUs, particularly at the edge.

Important STP timers are max age, hello time, and forward delay. The default values for the timers are:

- Hello time – 2 seconds
- Forward delay – 15 seconds
- Max age – 20 seconds

The operational values from these timers are derived from the root bridge. The current IEEE standard for Spanning Tree fixes hello time at 2 seconds. The Enterasys switches covered in this document do not enforce this restriction to allow existing configurations to remain compatible. It is not recommended that a value other than 2 seconds be used. Other values may not interact well with other non-variable protocol times such as edgeDelayWhile or mDelayWhile. The max age timer may be adjusted to change the network diameter. Take care to consider that failure in the network may cause the topology to “unravel” causing the diameter to become larger than anticipated. An insufficient value could cause devices near or at the edge of the network to become unreachable. For example, in a ring topology of 10 bridges, no bridge is more than 5 hops from the root. A max age that accounts for 6 hops would be sufficient. A failure of ports immediately interconnecting a bridge with the root would break the ring topology and change the furthest hop from the root from 5 to 9. Any bridges beyond the configured network diameter of 6 would cause the Spanning Tree topology not to converge.

## MSTP Operation

MSTP makes it possible for VLAN switching devices to use multiple Spanning Trees, allowing traffic belonging to different VLANs to flow over potentially different paths within the LAN. It builds upon the advancements of RSTP with its decreased time for network re-spans. MSTP's principle objective is to increase bandwidth utilization by allowing:

- Frames assigned to different VLANs to follow different data routes
- Ports to block for some Spanning Trees and forward for others
- Every inter-switch link in the topology to be forwarding for at least one Spanning Tree

MSTP is the default Spanning Tree mode on all Extreme Networks switch devices.

## Common and Internal Spanning Tree (CIST)

MSTP uses all Spanning Tree region information to create a single Common and Internal Spanning Tree (CIST) that represents the connectivity of the entire network. This is equivalent to the single Spanning Tree used for STP and RSTP.

The MSTP enabled network may contain any combination of Single Spanning Tree (SST) regions and Multiple Spanning Tree (MST) regions. A typical network may contain multiple MST regions as well as separate LAN segments running legacy STP and RSTP Spanning Tree protocols. The CIST contains a root bridge, which is the root of the Spanning Tree for the network. The CIST root may be, but is not necessarily, located inside an MST region. Each MST region contains a CIST regional root which may be the CIST root if the CIST root is internal to the region. If the CIST root is external to the region, the CIST regional root provides the connectivity to the CIST root. Bridges in an MSTP topology compare their received BPDUs to calculate their shortest path to the CIST root, CIST regional root, and MSTI regional root.




Ideally, there should be one all-encompassing region. This is not always possible, for example, when non-MSTP bridges exist such as those shown in [Figure 18-3](#) on page 5. From the outside, the region appears as a single Spanning Tree bridge which is part of the Common Spanning Tree (CST). A port which connects to a bridge not having the same MST configuration ID, or which is not running MSTP, forms part of the boundary of the region. The region attaches to the CST at the root port of the CIST regional root. All other region boundary ports which provide paths to the root are alternate ports and remain blocking until the topology changes, causing a new regional root port to be chosen. Ports which provide a path to the root for other bridges at the region boundary are designated ports. At boundary ports, port states for MSTIs follow the states of the CIST for the port.

## MST Region

An MST region is a group of devices that are configured together to form a logical region. The MST region presents itself to the rest of the network as a single switching device, which simplifies administration. Path cost is only incremented when traffic enters or leaves the region, regardless of the number of devices within the region. Each LAN can only be a member of one region. [Figure 18-7](#) shows that the MST region appears as a single switching device to devices 1 and 2, but really consists of three devices.

**Figure 18-7 Example of an MST Region**



For a switching device to be considered as part of an MST region, it must be administratively configured with the same configuration identifier information as all other devices in the MST region. The configuration identifier consists of four parts:

- Format Selector – One octet in length and is always 0. It cannot be administratively changed.
- Configuration Name – A user-assigned, case sensitive name given to the region. The maximum length of the name is 32 octets. A bridge's default configuration name is a character string corresponding to the bridge MAC address. This guarantees that the default behavior of a bridge is to not be part of an MST region.
- Revision Level – Two octets in length. The default value of 0 may be administratively changed.
- Configuration Digest – 16-octet HMAC-MD5 signature created from the configured VLAN Identification (VID)/Filtering Identification (FID) to Multiple Spanning Tree Instances (MSTI) mappings. All devices must have identical mappings to have identical configuration digests.

By default, each bridge is in its own MST region and has a default configuration name derived from the bridge MAC address. For example, if the bridge MAC address is **00-1f-45-9a-6c-b7**, the

default MSTP configuration name is “00:1f:45:9a:6c:b7”. When grouping two or more bridges into a single MST region, you must assign the same configuration name to each member of the region. MD5 digests are derived from a mapping of a Filtering Database ID (FID) to a Spanning Tree ID (SID), referred to as a FID-to-SID mapping (see [Multiple Spanning Tree Instances \(MSTI\)](#) on page 18-17 for more information). Since there is a small probability of different mappings resulting in the same digest, the addition of administratively assigned name and version configuration ID parameters guarantee the uniqueness of the region.

SIDs exist within an MST region, each having a separate topology. Within an MST region there always exists the Internal Spanning Tree (IST) which is SID 0. There are zero or more Multiple Spanning Tree Instances (MSTIs). Each MSTI corresponds to a set of VLANs. One or more VLANs may be mapped to an SID using a FID-to-SID mapping. The IST and each MSTI may have different root bridges. Port path costs and bridge priorities may be different for each port/instance. Each bridge port has a unique port state per instance. With proper configuration, redundant links may be utilized to their maximum extent by each forwarding for one or more instances. See [Configuring MSTP](#) on page 18-27 for more detail on how to do this.

## Multiple Spanning Tree Instances (MSTI)

Inside the MST region, a wholly contained set of topologies is maintained separate from the outside world. For example, MSTI 1 in MST region A has no correspondence to MSTI 1 in MST region B. The 7100-Series supports MST instances.

The Extreme Networks switch device by default maps VLAN IDs (VIDs) to Filtering IDs (FIDs) in a one-to-one correlation for bridges with the VLAN learning mode set to individual VLAN learning (IVL). VID to FID can also be mapped in a many-to-one correlation for bridges with the VLAN learning mode set to shared VLAN learning (SVL). The VLAN learning mode and shared VLAN learning VID to FID mapping are set by configuring VLAN constraint using the **set vlan constraint** command.

For example, in an IVL bridge, FID 3 may contain VID 3 and FID 4 may contain VID 4. In an SVL bridge, FID 3 may contain VID 3 and FID 4 may contain VIDs 4 and 5. Regardless of the type of VLAN learning taking place, one or more FIDs may be mapped to a Spanning Tree Instance (SID). The end result is a mapping of VID to SID. SID topologies may then be configured to provide a type of load balancing. Note that without further configuration, each SID will have the same topology as the IST. Typically, load balancing will be achieved by choosing different root bridges in the core for the different instances.

## Determining FID-to-SID Mappings

VLANs are mapped to MSTIs through a FID-to-SID mapping which is the key element in an MSTP configuration. Each VLAN is associated to a FID and is mapped to Spanning Tree IDs using their FID association. The mapping is performed by the **set spantree mstmap** command. This mapping is represented within the MST configuration digest described in the previous section and displayed in the following example. By default, every bridge will have a FID-to-SID mapping that equals VLAN FID 1/SID 0.

Use the **show spantree mstcfgid** command to determine MSTI configuration identifier information, and whether or not there is a misconfiguration due to non-matching configuration identifier components:

This example shows how to display MSTI configuration identifier information. In this case, this bridge belongs to “Region1”:

```
Extremenetworks->show spantree mstcfgid
MST Configuration Identifier:
  Format Selector:      0
  Configuration Name:  Region1
```

```
Revision Level:      88
Configuration Digest: 6d:d7:93:10:91:c9:69:ff:48:f2:ef:bf:cd:8b:cc:de
```

In order for other bridges to belong to Region1, all four elements of those bridges' configuration id output must match. The default value that must be changed for this to happen is the configuration name setting. Also, the MSTIs must be created and the FIDs mapped to them.

Use the **set spantree mstcfgid** command to change the configuration name from the default bridge MAC address value.

This example changes the default bridge configuration name to **Region1**:

```
Extremenetworks->set spantree mstcfgid cfgname Region1
```

For the configuration digest to match, the mapping of VIDs to SIDs must match. Use these commands to configure the SIDs, map the FIDs to the SIDs and display the VID-SID and FID-SID mappings:

```
Extremenetworks->set spantree msti sid 3 create
```

```
Extremenetworks->set spantree msti sid 4 create
```

```
Extremenetworks->set spantree mstmap 3 sid 3
```

```
Extremenetworks->set spantree mstmap 4 sid 4
```

```
Extremenetworks->show spantree mstlist
```

```
Configured Multiple Spanning Tree Instances:
```

```
  3      4
```

```
Extremenetworks->show spantree mstmap
```

```
Fid 3 is mapped to Sid 3
```

```
Fid 4 is mapped to Sid 4
```

```
Extremenetworks->show spantree vlanlist
```

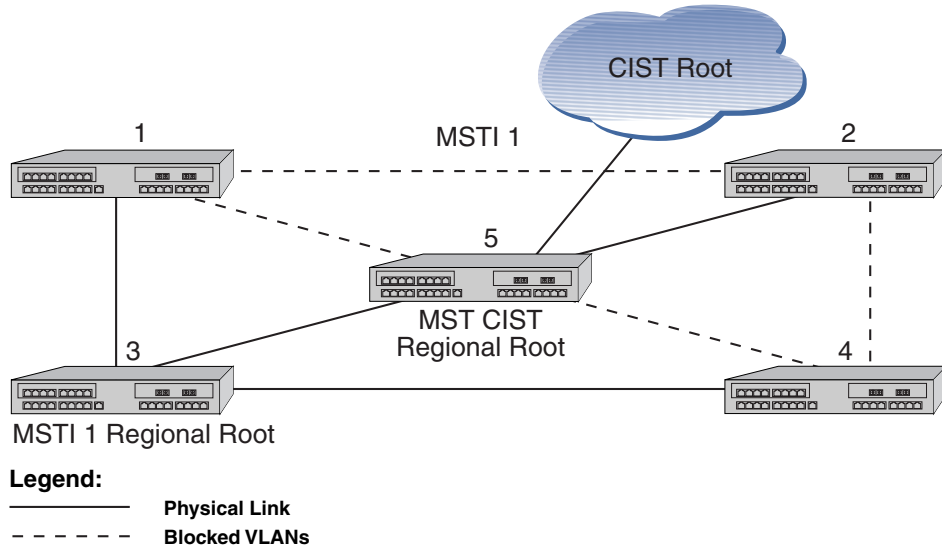
```
Vlan 3 is mapped to Sid 3
```

```
Vlan 4 is mapped to Sid 4
```

Since an MSTI is a separate Spanning Tree, each MSTI has its own root inside the MST region. [Figure 18-8](#) and [Figure 18-9](#) show two MSTIs in a single region. Switching device 3 is the root for MSTI 1, switching device 2 is the root for MSTI 2, and switching device 5 is the CIST regional root. Traffic for all the VLANs attached to an MSTI follow the MSTI's spanned topology.

Various options may be configured on a per-MSTI basis to allow for differing topologies between MSTIs. To reduce network complexity and processing overhead needed to maintain MSTIs, you should only create as many MSTIs as needed.

**Figure 18-8 MSTI 1 in a Region**



**Figure 18-9 MSTI 2 in the Same Region**

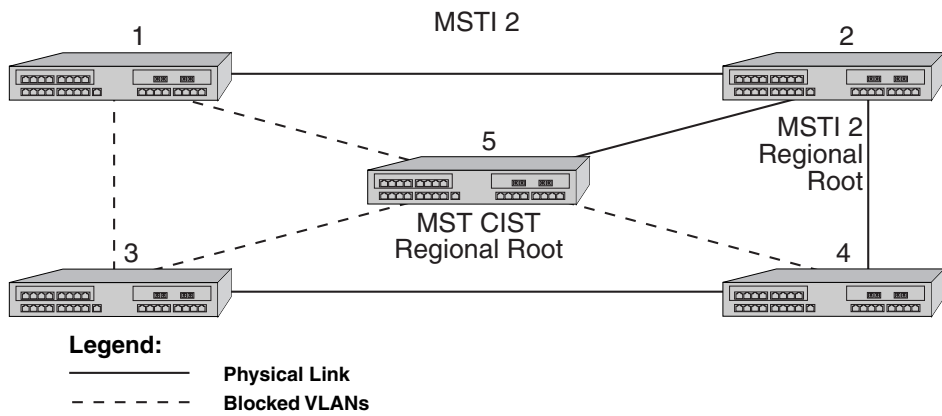



Figure 18-10 shows 3 regions with five MSTIs. Table 18-3 defines the characteristics of each MSTI. Ports connected to PCs from devices 1, 3, 9, and 11 will be automatically detected as edge ports. Devices 4 and 10 are the CIST regional roots. Each MSTI can be configured to forward and block various VLANs.

**Figure 18-10 Example of Multiple Regions and MSTIs**



**Table 18-3 MSTI Characteristics for Figure 18-10**

MSTI / Region	Characteristics
MSTI 1 in Region 1	Root is switching device 4, which is also the CIST regional root
MSTI 2 in Region 1	Root is switching device 5
MSTI 1 in Region 2	Root is switching device 7, which is also the CIST root
MSTI 1 in Region 3	Root is switching device 11
MSTI 2 in Region 3	Root is switching device 12
	Switching device 10 is the CIST regional root

## Multisource Detection

Multisource Detection is a feature that prevents network disruption due to excessive topology changes caused by a full duplex port transmitting multiple BPDUs with different source MAC addresses, and hence different BPDU information. When a port is point-to-point, the received priority information comes from the most recently received BPDU. When a port is non-point-to-point, the received information reflects the best priority information out of all the received BPDUs.

Typical scenarios for multisource detection are when a switch is connected to a device which:

- Has been improperly configured to forward received BPDUs out other ports
- Or has been configured to not run the Spanning Tree protocol and treats BPDUs as multicast packets by transmitting them out all other forwarding ports.

In these situations, the connected port is effectively acting as a shared media device. Shared media is detected using the duplex setting. Since the port is full duplex it treats the connection as point-to-point.

One way of preventing the disruption of this situation is to configure the receiving port's adminpoint value to false. This causes the operpoint value to always be false and to be treated as non-point-to-point. Multisource Detection, which is always enabled, recognizes the multiple source MAC addresses and automatically sets the operpoint value to false when the adminpoint value is auto. The port is constantly monitored. If the situation is resolved, as determined by

receiving a unique address for a sufficient amount of time, the operpoint value will be restored to true.

A syslog message is issued when multiple source addresses are detected:

```
Receive Event: Multiple BPDU sources received on Port = ge.1.1 BPDU Source MAC =
00:00:00:01:02:03 Prior Source MAC = 00:00:00:0f:0e:0d
```



**Note:** When loop protect is configured for the port, if multisource detection is triggered, the port will go to the listening state and no longer be part of the active topology. Loop protect does not operate on shared media ports.

## Configuring STP and RSTP



**Caution:** Spanning Tree configuration should be performed only by personnel who are very knowledgeable about Spanning Trees and the configuration of the Spanning Tree Algorithms. Otherwise, the proper operation of the network could be at risk.

For information about...	Refer to page...
<a href="#">Reviewing and Enabling Spanning Tree</a>	<a href="#">18-21</a>
<a href="#">Adjusting Spanning Tree Parameters</a>	<a href="#">18-22</a>
<a href="#">Enabling the Backup Root Function</a>	<a href="#">18-25</a>
<a href="#">Adjusting RSTP Parameters</a>	<a href="#">18-26</a>

## Reviewing and Enabling Spanning Tree

By default, Spanning Tree is enabled globally on Extreme Networks switch devices and enabled on all ports. On all switching devices, the default Spanning Tree version is set to MSTP (802.1s) mode. Since MSTP mode is fully compatible and interoperable with legacy STP and RSTP bridges, in most networks, this default should not be changed.

Use the following commands to review, re-enable, and reset the Spanning Tree mode.

- Review the current configuration on one or more SIDs, ports, or both:
 

```
show spantree stats [port port-string] [sid sid] [active]
```

 Specifying **active** will display information for port(s) that have received BPDUs since boot.
- If necessary, globally enable Spanning Tree:
 

```
set spantree stpmode ieee8021
```
- Review the status of Spanning Tree on one or more ports:
 

```
show spantree portadmin [port port-string]
```
- If necessary, re-enable Spanning Tree on one or more ports:
 

```
set spantree portadmin port-string enable
```

### Example

This example shows how to display the device's Spanning Tree configuration:

```
Extremenetworks->show spantree stats
SID - 1
Spanning tree mode - enabled
```

```

Designated Root          - 00-e0-63-6c-9b-6d
Designated Root Priority - 0
Designated Root Cost     - 1
Designated Root Port     - ge.5.1
Root Max Age             - 20 sec
Root Hello Time          - 2 sec
Root Forward Delay       - 15 sec
Bridge ID MAC Address    - 00-e0-63-9d-b5-87
Bridge priority          - 32768
Bridge Max Age           - 20 sec
Bridge Hello Time        - 2 sec
Bridge Forward Delay     - 15 sec
Topology Change Count    - 6539
Time Since Top Change    - 00 days 00:00:00

```



**Note:** By default, Spanning Tree is enabled both globally and on all ports.

## Adjusting Spanning Tree Parameters

You may need to adjust certain Spanning Tree parameters if the default values are not suitable for your bridge configuration. Parameters affecting the entire Spanning Tree are configured with variations of the global bridge configuration commands. Interface-specific parameters are configured with variations of the Spanning Tree port configuration commands. Default settings are listed in [Table 18-4](#):

**Table 18-4 Spanning Tree Port Default Settings**

Setting	Default Value
Bridge priority mode	802.1t
Bridge priority	32768
Port priority	128
Port cost	0 (automatically calculated based on port speed)
Hello time (bridge and ports)	2 seconds
Bridge forward delay	15 seconds
Bridge maximum aging time	20 seconds

Use the commands in the following sections to adjust these defaults.



**Note:** Poorly chosen adjustments to these parameters can have a negative impact on network performance. Please refer to the IEEE 802.1D specification for guidance.

### Setting Bridge Priority Mode and Priority

Bridge priority mode affects the range of priority values used to determine which device is selected as the Spanning Tree root. By default, switching devices are set to 802.1t mode as described in [Updated 802.1t](#) on page 18-8.

Use this command to set the bridge priority mode:

```
set spantree bridgepriortymode 802.1t | 802.1d
```

In addition to setting priority mode, you can globally configure the priority of an individual bridge. When two bridges tie for position as the root bridge, this setting affects the likelihood that a bridge will be selected. The lower the bridge's priority, the more likely the bridge will be selected as the root bridge.

Use this command to set the bridge priority:

```
set spantree priority priority [sid]
```

Valid *priority* values are:

- For 802.1t priority mode: **0–61440** (in increments of 4096), with 0 indicating high priority and 61440 low priority. Values will automatically be rounded up or down, depending on the 802.1t value to which the entered value is closest.
- For 802.1D priority mode: **0–65535** (in increments of 1), with 0 indicating high priority and 65535 low priority.

Valid *sid* values are **0–4094**. If not specified, SID 0 will be assumed.

## Setting a Port Priority

You can set a Spanning Tree port priority. Port priority is used to break a tie when choosing the root port for a bridge, in a case where the choice is between ports connected to the same bridge. The port with the lowest value will be elected.

Use this command to set a port priority:

```
set spantree portpri port-string priority [sid sid]
```

Valid *priority* values are **0–240** (in increments of 16) with 0 indicating high priority.

Valid *sid* values are **0–4094**. If not specified, SID 0 will be assumed.

## Assigning Port Costs

Each interface has a Spanning Tree port cost associated with it, which helps to determine the quickest path between the root bridge and a specified destination. By convention, the higher the port speed, the lower the port cost. By default, this value is set to 0, which forces the port to recalculate Spanning Tree port cost based on the speed of the port and whether or not legacy (802.1D) path cost is enabled.

Use this command to assign different Spanning Tree port costs:

```
set spantree adminpathcost port-string cost [sid sid]
```

Valid *cost* values are:

- **0–65535** if legacy path cost is enabled.
- **0–200000000** if legacy path cost is disabled.

Valid *sid* values are **0–4094**. If not specified, SID 0 will be assumed.



**Notes:** Please refer to the IEEE 802.1D specification for guidance in setting appropriate cost values for your port speeds.

By default, legacy path cost is disabled. Enabling the device to calculate legacy path costs affects the range of valid values that can be administratively assigned.

To check the status of legacy path cost, use **show spantree legacypathcost**.

To disable legacy path cost, if necessary use **set spantree legacypathcost disable**.



## Adjusting Bridge Protocol Data Unit (BPDU) Intervals

Use the commands in this section to adjust default BPDU interval values.

**Table 18-5 BPDU Interval Defaults**

BPDU Interval	Default Value
Hello time (bridge and ports)	2 seconds
Forward delay	15 seconds
Maximum age time	20 seconds

### Adjusting the Bridge Hello Time



**Caution:** Poorly chosen adjustments to bridge and port hello time parameters can have a negative impact on network performance. It is recommended that you do not change these parameters unless you are familiar with Spanning Tree configuration and have determined that adjustments are necessary. Please refer to the IEEE 802.1D specification for guidance.

Hello time is the interval, in seconds, at which the bridge or individual ports send BPDU messages. By default, bridge hello mode is enabled, meaning the device uses a single bridge administrative hello time.

Adjust the bridge hello time as follows:

1. Check the status of bridge hello mode:
 

```
show spantree bridgehellomode
```
2. If necessary, re-enable bridge hello mode:
 

```
set spantree bridgehellomode enable
```
3. Set a new hello time interval:
 

```
set spantree hello interval
```

Valid *interval* values are 1–10.

### Adjusting Port Hello Times

You can set the device to use per-port administrative hello times by disabling bridge hello mode and adjusting the hello time interval for one or more ports as follows:

1. Check the status of bridge hello mode:
 

```
show spantree bridgehellomode
```
2. If necessary, disable bridge hello mode:
 

```
set spantree bridgehellomode disable
```
3. Set a new hello time interval for one or more ports:
 

```
set spantree porthello port-string interval
```

Valid *interval* values are 10–100

### Adjusting the Forward Delay Interval

When rapid transitioning is not possible, forward delay is used to synchronize BPDU forwarding. The forward delay interval is the amount of time spent listening for topology change information after an interface has been activated for bridging and before forwarding actually begins. This delay is required because every device must receive information about topology changes before it

starts to forward frames. In addition, each port needs time to listen for conflicting information that would make it return to a blocking state. Otherwise, temporary data loops might result.

Use this command to adjust the forward delay interval setting:

```
set spantree fwddelay delay
```

Valid *delay* values are 4–30.

## Defining the Maximum Age Time

If a bridge does not hear BPDUs from the root bridge within the interval (number of seconds) specified as maximum age time, it assumes that the network has changed and recomputes the Spanning Tree topology. By adjusting this value, you can configure support for a maximum diameter from the STP root of up to 40 bridges. By default, Extreme Networks switching devices are set with a maximum age time of 20 seconds, supporting a 20-bridge span from the root bridge.

Use this command to adjust the maximum age setting:

```
set spantree maxage agingtime
```

Valid *agingtime* values are 6–40 (seconds).

## Setting the Maximum Configurable STPs

By default, Multiple Spanning Tree mode is globally enabled on Extreme Networks switching devices and one Spanning Tree is configured as Spanning Tree ID (SID) 0. As described in [“Maximum SID Capacities”](#) on page 18-6, devices support different numbers of Spanning Tree instances (including SID 0), depending on their model type and memory installed. SID values are from 1 to 4094.

The 7100-Series allows you to set the maximum number of user configured Spanning Trees allowed on the device:

```
set spantree maxconfigurablestps numstps
```

Valid *numstps* values are

## Setting Restricted TCN and Restricted Role

As described in [Restricted Topology Change Notification \(TCN\)](#) on page 18-8, Restricted Topology Change Notice (TCN) allows or disallows TCN propagation on specified ports.

The 7100-Series device allows you to restrict the propagation of TCNs on the specified port by setting Restricted TCN to true:

```
set spantree restrictedtcn port-string {true | false}
```

By default, Restricted TCN is set to false, allowing the propagation of TCNs on the specified port.

As described in [Restricted Role](#) on page 18-9, Restricted Role allows or disallows the root role on specified ports.

The 7100-Series device allows you to restrict the root role on the specified port by setting Restricted Role to true:

```
set spantree restrictedrole port-string {true | false}
```

By default, Restricted Role is set to false, allowing the root role on the specified port.

## Enabling the Backup Root Function

Disabled by default, the backup root function works only when the backup root-enabled bridge is directly connected to the root bridge. The backup root function prevents stale Spanning Tree

information from circulating throughout the network in the event that the link between the root bridge and the backup root-enabled bridge is lost. If this happens, the backup root will dynamically lower its bridge priority relative to the existing root bridge's priority, causing it to immediately be selected as the new root bridge.

Use this command to enable the backup root function on an SID:

```
set spantree backuproot sid enable
```

When SNMP trap messaging is configured and the backup root function is enabled, a trap message will be generated when the backup becomes the new root of the network.

## Adjusting RSTP Parameters

Since rapid link reconfiguration can happen only on a point-to-point link or an edge port (a port that is known to be on the edge of a bridged LAN), in some cases you may want to define them administratively. However, since edge port and point-to-point links are automatically detected on Extreme Networks switching devices, in most cases you will not need to change these default port designations.

### Defining Point-to-Point Links

By default, the administrative point-to-point status is set to auto on all Spanning Tree ports, allowing the Extreme Networks firmware to determine each port's point-to-point status. In most cases, this setting will not need to be changed and will provide optimal RSTP functionality. You can, however, use the following commands to review and, if necessary, change the point-to-point status of a Spanning Tree link.

Review and define the point-to-point status of an RSTP link as follows:

1. Display the point-to-point operating status of a LAN segment attached to a port:

```
show spantree operpoint [port port-string]
```

A status of "true" indicates the LAN segment is operating as a point-to-point link.

A status of "false" indicates it is not.

If *port-string* is not specified, point-to-point operating status will be displayed for all Spanning Tree ports.

2. Display the point-to-point administrative status of a LAN segment attached to a port:

```
show spantree adminpoint [port port-string]
```

A status of "true" indicates the port is administratively set to be considered point-to-point.

A status of "false" indicates the port is administratively set to be considered non point-to-point.

A status of "auto" (the default setting) indicates that the firmware is allowed to determine the port's point-to-point status.

If *port-string* is not specified, point-to-point administrative status will be displayed for all Spanning Tree ports.

3. If necessary, change the point-to-point administrative status of a LAN segment attached to a port:

```
set spantree adminpoint port-string {auto | true | false}
```

## Defining Edge Port Status

By default, edge port status is disabled on all ports. When enabled, this indicates that a port is on the edge of a bridged LAN. You can use the following commands to review and, if necessary, change the edge port detection status on the device and the edge port status of Spanning Tree ports.

Review and define edge port status as follows:

1. Display the status of edge port detection:

```
show spantree autoedge
```

2. If desired, enable edge port detection:

```
set spantree autoedge enable
```

3. Display the edge port operating status of one or more port(s):

```
show spantree operedge [port port-string]
```

A status of "true" or "Edge-Port" indicates the port is operating as an edge port.

A status of "false" or "Non-Edge-Port" indicates it is not.

If *port-string* is not specified, edge port status will be displayed for all Spanning Tree ports.

4. Display the edge port administrative status of one or more port(s):

```
show spantree adminedge [port port-string]
```

A status of "true" or "Edge-Port" indicates the port is administratively set to be considered an edge port.

A status of "false" or "Non-Edge-Port" indicates the port is administratively set to be considered a non edge port.

If *port-string* is not specified, edge port administrative status will be displayed for all Spanning Tree ports.

5. If necessary, change the edge port administrative status of one or more port(s):

```
set spantree adminedge port-string true
```

## Configuring MSTP

In order for MSTP to provide multiple forwarding paths, the following must happen:

- The configuration identifier must match on all bridges within the region.
- All bridges must be within the same region.
- All bridges must be connected to MSTP-aware bridges. (They can be connected using a shared media such as a repeater provided that a single Spanning Tree device does not reside on that LAN).



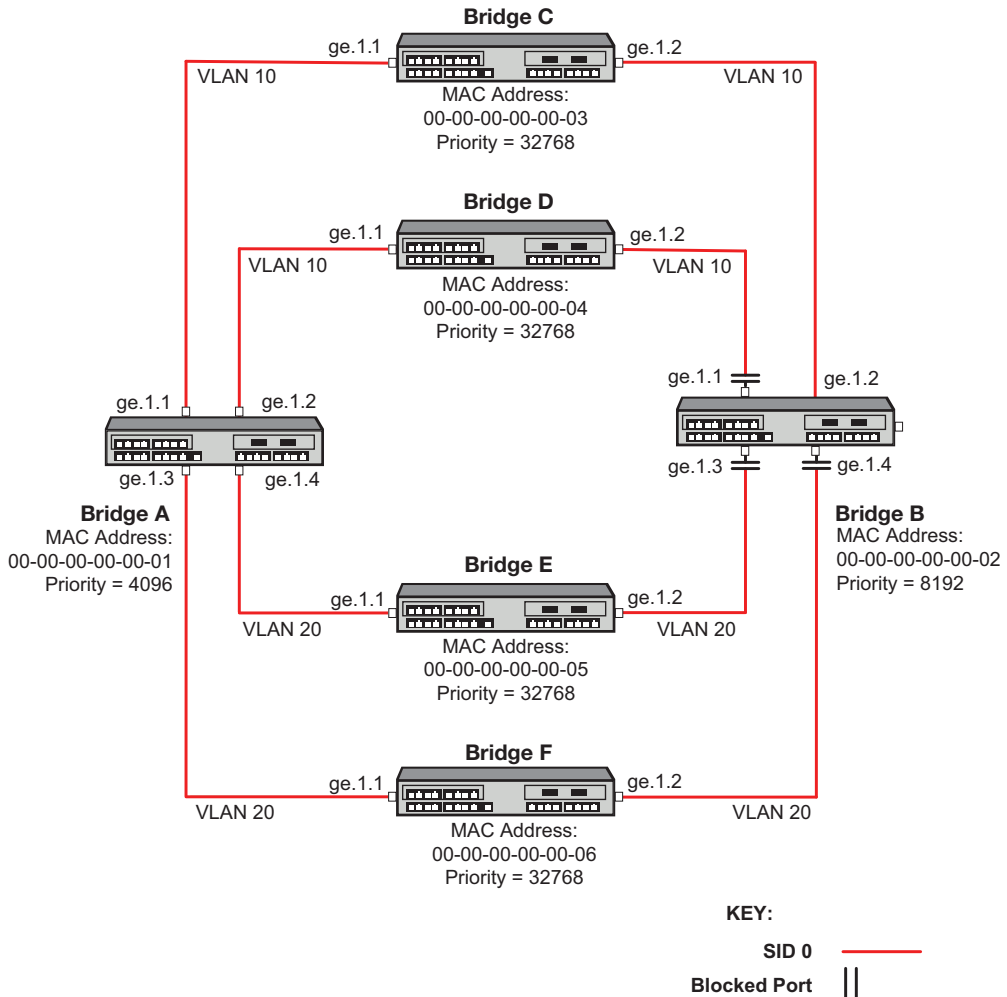
**Note:** A single Spanning Tree device between two MSTP bridges will terminate the ability to have multiple forwarding paths.

For information about...	Refer to page...
<a href="#">Example 1: Configuring MSTP for Traffic Segregation</a>	18-28
<a href="#">Example 2: Configuring MSTP for Maximum Bandwidth Utilization</a>	18-30
<a href="#">Adjusting MSTP Parameters</a>	18-31

### Example 1: Configuring MSTP for Traffic Segregation

This example illustrates the use of MSTP for traffic segregation by VLAN and SID. Bridges A, B, C and D participate in VLAN 10. Bridges A, B, E and F participate in VLAN 20. Figure 18-11 shows the problem that arises when using a single Spanning Tree configuration for traffic segregation with redundancy.

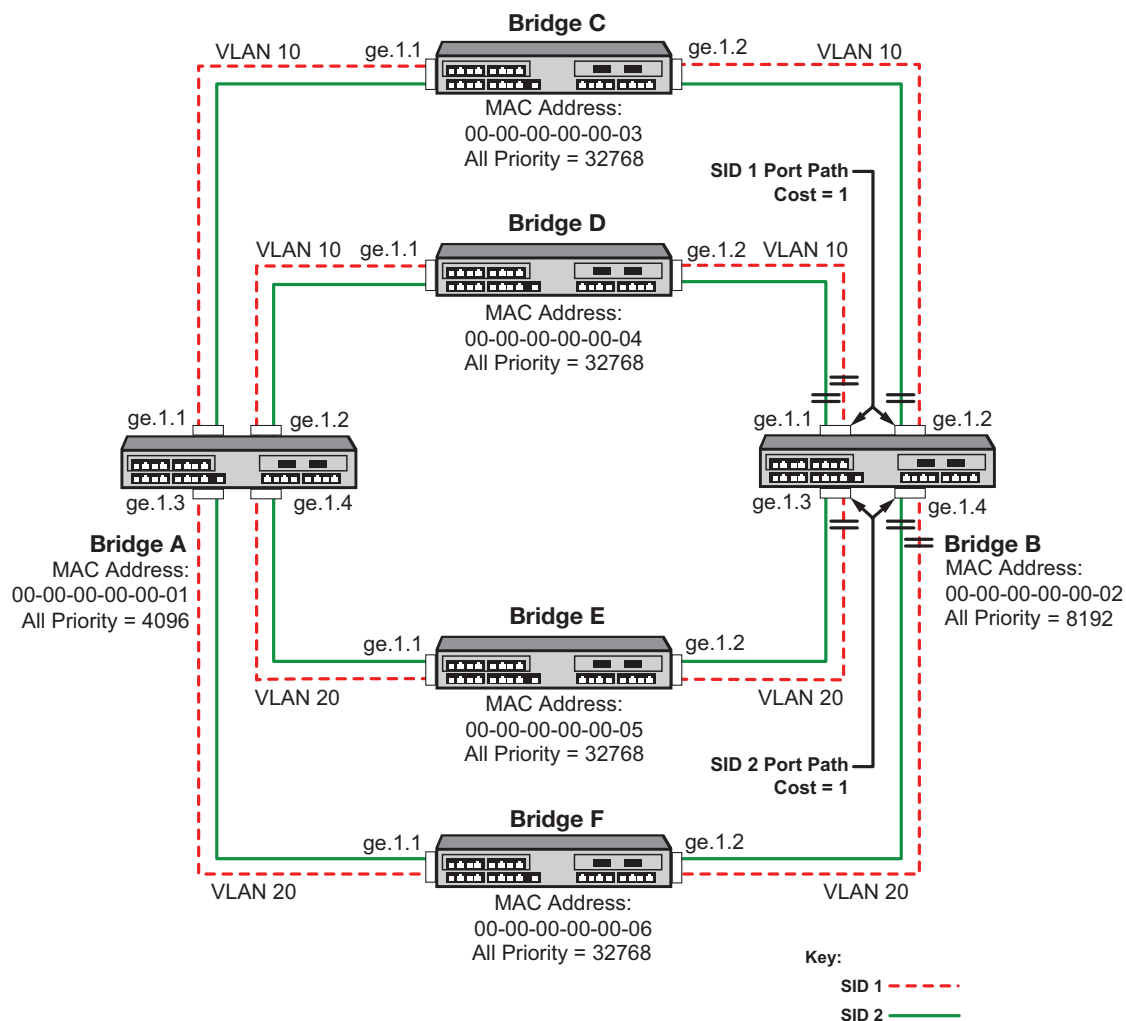
**Figure 18-11 Traffic Segregation in a Single STP Network Configuration**



In a single Spanning Tree configuration a bridge can only have one port forwarding towards the root for all traffic. Bridge A has the lowest priority and is the root. Bridge B forwards traffic towards the root on port ge.1.2. All other ports are blocked. For this configuration, Bridge B will not have any active links forwarding for VLAN 20.

Figure 18-12 shows the solution using MSTP. By configuring separate Spanning Tree instances to overlay the two VLAN topologies, Bridge B port ge.1.2 forwards on VLAN 10 for SID 1 and port ge.1.3 forwards on VLAN 20 for SID 2.

Figure 18-12 Traffic Segregation in an MSTP Network Configuration



To configure the traffic segregation MSTP example on all bridges:

- Configure the MST configuration ID with the same name  
`set spantree mstcfgid cfgname name`
- Create SIDs 1 and 2  
`set spantree msti sid sid create`
- Create the FID to SID mappings VLAN 10 to SID 1 and VLAN 20 to SID 2  
`set spantree mstmap vlan-id sid sid`

To configure Bridge A as root, set the priority to 4096 for both SID 1 and SID 2.

`set spantree priority priority sid`

To configure Bridge B as the backup should Bridge A fail:


- Set the Spanning Tree priority to 8192 for both SID 1 and SID 2  
`set spantree priority priority sid`
- Set the admin path cost on ports ge.1.1-2 to 1 for SID 1
- Set the admin path cost on ports ge.1.3-4 to 1 for SID 2

```
set spantree adminpathcost port-id cost sid
```

## Example 2: Configuring MSTP for Maximum Bandwidth Utilization


This example illustrates the use of MSTP for maximum bandwidth utilization. Maximum bandwidth utilization takes place when all bridges participate on all VLANs. [Figure 18-13](#) shows that with a single Spanning Tree configuration, only a single link towards the root forwards on a bridge. The alternate ports are blocking.

**Figure 18-13 Maximum Bandwidth Utilization in a Single STP Network Configuration**



In [Figure 18-13](#), Bridge A is the root of the Spanning Tree because it has the lowest priority. Bridge D port ge.1.2 forwards traffic to Bridge A. Bridge D port ge.1.1 is blocking. Bridge C port ge.1.1 forwards traffic to Bridge A. Bridge C port ge.1.2 is blocking. This single Spanning Tree configuration prevents maximum bandwidth utilization for this network.

[Figure 18-14](#) on page 18-31 shows that with an MSTP configuration each link can be forwarding for some VLAN and each VLAN has a path to the root bridge.

**Figure 18-14 Maximum Bandwidth Utilization in an MSTP Network Configuration**

To configure the MSTP maximum bandwidth utilization example on all bridges:

- Create VLANs 10 and 20
 

```
set vlan create vlan-id
```
- Configure the MST configuration ID with the same name
 

```
set spanntree mstcfigid cfname name
```
- Create SIDs 86 and 99
 

```
set spanntree msti sid sid create
```
- Create the FID to SID mappings VLAN 10 to SID 86 and VLAN 20 to SID 99
 

```
set spanntree mstmap vlan-id sid sid
```

Additionally, the root of each SID is chosen to be in a different bridge. This will spread out the traffic. The bridges on the next level down have a link to each of the root bridges.

To configure Bridge A as root for SID 86, set the priority to 4096 for SID 86.

```
set spanntree priority priority sid
```

To configure Bridge B as the root for SID 99, set the priority to 4096 for SID 99.

## Adjusting MSTP Parameters

You may need to adjust certain Spanning Tree parameters if the default values are not suitable for your bridge configuration. Refer back to [Adjusting Spanning Tree Parameters](#) on page 18-22 and [Adjusting RSTP Parameters](#) on page 18-26 for information on adjusting Spanning Tree defaults. Changes made to global and port-related Spanning Tree defaults will take affect if the device is running in STP, RSTP, or MSTP.



## Monitoring MSTP

Use the commands in [Table 18-6](#) to monitor MSTP statistics and configurations. You can also use the show commands described in [Reviewing and Enabling Spanning Tree](#) on page 18-21 to review information related to all Spanning Tree protocol activity.

**Table 18-6 Commands for Monitoring MSTP**

Task	Command
Verify that MSTP is running on the device.	<code>show spantree version</code>
Display the maximum configurable MSTIs allowed on the device.	<code>show spantree maxconfigurablesteps</code>
Display a list of MSTIs configured on the device.	<code>show spantree mstlist</code>
Display the mapping of one or more filtering database IDs (FIDs) to Spanning Trees. Since VLANs are mapped to FIDs, this shows to which SID a VLAN is mapped.	<code>show spantree mstmap [fid fid]</code>
Display the Spanning Tree ID(s) assigned to one or more VLANs.	<code>show spantree vlanlist [vlan-list]</code>
Display MST configuration identifier elements, including format selector, configuration name, revision level, and configuration digest.	<code>show spantree mstcfgid</code>
Display protocol-specific MSTP counter information.	<code>show spantree debug [port port-string] [sid sid] [active]</code>

## Understanding and Configuring SpanGuard

For information about...	Refer to page...
<a href="#">What Is SpanGuard?</a>	<a href="#">18-32</a>
<a href="#">How Does It Operate?</a>	<a href="#">18-34</a>
<a href="#">Configuring SpanGuard</a>	<a href="#">18-33</a>

### What Is SpanGuard?

As described previously in the overview of [SpanGuard](#) on page 18-7, this feature enables Extreme Networks switching devices to detect unauthorized bridges in your network, resolving the threat of repeated topology change notifications or new root bridge announcements causing a Denial of Service (DoS) condition. It prevents Spanning Tree respsans that can occur when BPDUs are received on user ports and notifies you (network management) they were attempted.

If a SpanGuard enabled port receives a BPDU, it becomes locked and transitions to the blocking state. It will only transition out of the blocking state after a globally specified time or when it is manually unlocked.

By default, SpanGuard is globally disabled and must be globally enabled to operate on all user ports. For configuration information, refer to [Configuring SpanGuard](#) on page 18-33.

## How Does It Operate?

SpanGuard helps protect against Spanning Tree Denial of Service (DoS) SpanGuard attacks as well as unintentional or unauthorized connected bridges, by intercepting received BPDUs on configured ports and locking these ports so they do not process any received packets.

When enabled, reception of a BPDU on a port that is administratively configured as a Spanning Tree edge port (`admedge = True`) will cause the port to become locked and the state set to blocking. When this condition is met, packets received on that port will not be processed for a specified timeout period. The port will become unlocked when:

- the timeout expires,
- the port is manually unlocked,
- the port is no longer administratively configured as `admedge = True`, or
- the SpanGuard function is disabled.

The port will become locked again if it receives another offending BPDU after the timeout expires or it is manually unlocked.

In the event of a DoS attack with SpanGuard enabled and configured, no Spanning Tree topology changes or topology reconfigurations will be seen in your network. The state of your Spanning Tree will be completely unaffected by the reception of any spoofed BPDUs, regardless of the BPDU type, rate received or duration of the attack.

By default, when SNMP and SpanGuard are enabled, a trap message will be generated when SpanGuard detects that an unauthorized port has tried to join a Spanning Tree.

## Configuring SpanGuard

Use the following commands to configure device ports for SpanGuard, to enable the SpanGuard function, and to review SpanGuard status on the device.

### Reviewing and Setting Edge Port Status



**Note:** To use the SpanGuard function, you must know which ports are connected between switching devices as ISLs (inter-switch links). Also, you must configure edge port status (`admedge = true` or `false`) on the entire switch, as described in “[Defining Edge Port Status](#)” on page 18-27, before SpanGuard will work properly.

Review and set edge port status as follows:

1. Use the show commands described in “[Defining Edge Port Status](#)” on page 18-27 to determine edge port administrative status on the device.
2. Set edge port administrative status to `false` on all known ISLs.
3. Set edge port administrative status to `true` on any remaining ports where SpanGuard protection is desired. This indicates to SpanGuard that these ports are not expecting to receive any BPDUs. If these ports do receive BPDUs, they will become locked.

### Enabling and Adjusting SpanGuard

Use this command to enable SpanGuard on the device:

```
set spantree spanguard enable
```

Use this command to adjust the SpanGuard timeout value. This sets the length of time that a SpanGuard-affected port will remain locked:

```
set spantree spanguardtimeout timeout
```

Valid values are 0–65535 seconds. Default is 300 seconds. Setting the value to 0 will set the timeout to forever.

Use this command to manually unlock a port that was locked by the SpanGuard function. This overrides the specified timeout variable:

```
set spantree spanguardlock port-string
```

## Monitoring SpanGuard Status and Settings

Use the commands in [Table 18-7](#) to review SpanGuard status and settings.

**Table 18-7 Commands for Monitoring SpanGuard**

Task	Command
Display the status of SpanGuard on the device.	<code>show spantree spanguard</code>
Display the status of the SpanGuard lock function on one or more ports.	<code>show spantree spanguardlock [port <i>port-string</i>]</code>
Display the SpanGuard timeout setting.	<code>show spantree spanguardtimeout</code>
Display the status of the SpanGuard trap function.	<code>show spantree spanguardtrapeenable</code>

## Understanding and Configuring Loop Protect

For information about...	Refer to page...
<a href="#">What Is Loop Protect?</a>	<a href="#">18-34</a>
<a href="#">How Does It Operate?</a>	<a href="#">18-34</a>
<a href="#">Configuring Loop Protect</a>	<a href="#">18-37</a>

### What Is Loop Protect?

As described previously in the overview of [Loop Protect](#) on page 8, this feature prevents or short circuits loop formation in your network. It does this by requiring ports to receive type 2 BPDUs (RSTP/MSTP) on point-to-point inter-switch links (ISLs) before their states are allowed to become forwarding. Further, if a BPDU timeout occurs on a port, its state becomes non-forwarding until a BPDU is received.

In this way, both upstream and downstream facing ports are protected. When a root or alternate port loses its path to the root bridge due to a message age expiration, it takes on the role of designated port and will not forward traffic until a BPDU is received.

When a port is intended to be the designated port in an ISL, it constantly proposes and will not forward until a BPDU is received. This protects against misconfiguration and protocol failure by the connected bridge.

### How Does It Operate?

Loop Protect operates as a per port, per MST instance feature and should be set on ISLs. It comprises several related functions, including:

- Controlling port forwarding state based on reception of agreement BPDUs
- Controlling port forwarding state based on reception of disputed BPDUs

- Communicating port non-forwarding status through traps and syslog messages
- Disabling a port based on frequency of failure events

## Port Modes and Event Triggers

Ports work in two Loop Protect operational modes. If the port is configured so that it is connected to a switching device known to implement Loop Protect, it uses full functional (enhanced) mode. Otherwise, it operates in limited functional (standard) mode.

Connection to a Loop Protect switching device guarantees that the alternate agreement mechanism is implemented and, therefore, the designated port can rely on receiving a response to its proposal regardless of the role of the connected port. This has two important implications. First, the designated port connected to a non-root port may transition to forwarding. Second, there is no ambiguity when a timeout happens; a Loop Protect event has occurred.

In full mode, when a type 2 BPDU is received and the port is designated and point-to-point, the timer is set to 3 times hello time. Limited mode adds a further requirement that the flags field in the BPDU indicates a root role. If the port is a boundary port, the MSTIs for that port follow the CIST (for example if the MSTI port timers are set according to the CIST port timer). If the port is internal to the region, the MSTI port timers are set independently using the particular MSTI message.

Loop Protect initializes the MSTI timer to zero and does not allow the designated port to transition from listening to learning until the timer becomes non-zero. If the port is not designated, the timer does not apply. Its state is controlled through normal protocol behavior.

A disputed BPDU is one in which the flags field indicates a designated role, a learning state, and the priority vector is worse than that already held by the port. If a disputed BPDU is received, the port is forced to the listening state.

Message age expiration and the expiration of the Loop Protect timer are both events for which Loop Protect generates a notice level syslog message. You can also configure traps to report these events, as well as a syslog message and trap for disputed BPDUs.

In addition, you can configure Loop Protect to force the locking of an SID/port when one or more events occur. When the configured number of events happen within a given window of time, the port will be forced into blocking and held there until you manually unlock it.

## Example: Basic Loop Protect Configuration

The following sample configuration shows how Loop Protect functions in a basic Spanning Tree topology.

In the example in [Figure 18-15](#), Switch 1 is the root bridge with BPDUs being sent to both Switch 2 and 3. (Designated ports are labeled D and root ports are labeled R.) Switch 3 has placed the port that connects to Switch 2 in a blocking state.

**Figure 18-15 Basic Loop Protect Scenario**




Figure 18-16 shows that, without Loop Protect, a failure could be as simple as someone accidentally disabling Spanning Tree on the port between Switch 2 and 3. Switch 3's blocking port eventually transitions to a forwarding state which leads to a looped condition.

**Figure 18-16 Spanning Tree Without Loop Protect**





Figure 18-17 shows that, with Loop Protect enabled, Switch 3 will not go to a forwarding state until it has received a BPDU from Switch 2.

**Figure 18-17 Spanning Tree with Loop Protect**



## Configuring Loop Protect

For information about...	Refer to page...
<a href="#">Enabling or Disabling Loop Protect</a>	<a href="#">18-37</a>
<a href="#">Specifying Loop Protect Partners</a>	<a href="#">18-37</a>
<a href="#">Setting the Loop Protect Event Threshold and Window</a>	<a href="#">18-37</a>
<a href="#">Enabling or Disabling Loop Protect Event Notifications</a>	<a href="#">18-38</a>
<a href="#">Setting the Disputed BPDU Threshold</a>	<a href="#">18-38</a>
<a href="#">Monitoring Loop Protect Status and Settings</a>	<a href="#">18-38</a>

### Enabling or Disabling Loop Protect

By default, Loop Protect is disabled on all ports. Use this command to enable (or, if desired, disable) the feature on one or more ports:

```
set spantree lp port-string {enable | disable} [sid sid]
```

If no SID is specified, SID 0 is assumed.

This command takes precedence over per port STP enable/disable state (portAdmin). Normally, portAdmin disabled would cause a port to go immediately to forwarding. If Loop Protect is enabled, that port should go to listening and remain there.



**Note:** The Loop Protect enable/disable settings for an MSTI port should match those for the CIST port.

### Specifying Loop Protect Partners

By default, each port is not set as a Loop Protect capable partner. If the port is set as a Loop Protect capable partner (true), the full functionality of the Loop Protect feature is used. If the value is false, then there is some ambiguity as to whether an Active Partner timeout is due to a loop protection event or is a normal situation due to the fact that the partner port does not transmit Alternate Agreement BPDUs. Therefore, a conservative approach is taken in that designated ports will not be allowed to forward unless receiving agreements from a port with root role. This type of timeout will not be considered a loop protection event. Loop protection is maintained by keeping the port from forwarding, but since this is not considered a loop event, it will not be factored into locking the port.

Use this command to set the Loop Protect partner state on one or more ports:

```
set spantree lpcapablepartner port-string {true | false}
```

### Setting the Loop Protect Event Threshold and Window

The Loop Protect event threshold is a global integer variable that provides protection in the case of intermittent failures. The default value is 3. If the event counter reaches the threshold within a given period (the event window), the port for the given SID becomes locked (that is, held indefinitely in the blocking state). If the threshold is 0, the ports are never locked.

Use this command to set the Loop Protect event threshold:

```
set spantree lpthreshold value
```

The Loop Protect window is a timer value, in seconds, that defines a period during which Loop Protect events are counted. The default value is 180 seconds. If the timer is set to 0, the event counter is not reset until the Loop Protect event threshold is reached.

Use this command to set the Loop Protect event window value in seconds:

```
set spantree lpwindow value
```

## Enabling or Disabling Loop Protect Event Notifications

Loop Protect traps are sent when a Loop Protect event occurs, that is, when a port goes to listening due to not receiving BPDUs. The trap indicates port, SID and loop protection status.

Use this command to enable or disable Loop Protect event notification. By default, this is disabled:

```
set spantree lptrapenable {enable | disable}
```

## Setting the Disputed BPDUs Threshold

A disputed BPDU is one in which the flags field indicates a designated role and a learning state, and the priority vector is worse than that already held by the port. If a disputed BPDU is received, the port is forced to the listening state. Refer to the 802.1Q-2005 standard, *IEEE Standard for Local and Metropolitan Area Networks – Virtual Bridged Local Area Networks*, for a full description of the dispute mechanism, which prevents looping in cases of one-way communication.

The disputed BPDU threshold is an integer variable that represents the number of disputed BPDUs that must be received on a given port and SID before a disputed BPDU trap is sent and a syslog message is issued. For example, if the threshold is 10, a trap is issued when 10, 20, 30 (and so on) disputed BPDUs have been received. The trap indicates port, SID and total Disputed BPDU count.

Use this command to set the disputed BPDU threshold:

```
set spantree disputedbpduthreshold value
```

Default value is 0, which means that traps are not sent.

## Monitoring Loop Protect Status and Settings

Use the commands in [Table 18-8](#) to monitor Loop Protect settings.

**Table 18-8** Commands for Monitoring Loop Protect

Task	Command
Display the Loop Protect status per port, per SID, or both.	<b>show spantree lp</b> [port <i>port-string</i> ] [sid <i>sid</i> ]
Display the Loop Protect lock status per port, per SID, or both. <b>Note:</b> A port can become locked if a configured number of Loop Protect events occur during the configured window of time. Once a port is forced into blocking (locked), it remains locked until manually unlocked with the <b>clear spantree llock</b> command.	<b>show spantree llock</b> [port <i>port-string</i> ] [sid <i>sid</i> ]
Display the Loop Protect capability of a link partner for one or more ports.	<b>show spantree lpcapablepartner</b> [port <i>port-string</i> ]
Display the reason for placing a port in a non-forwarding state due to an exceptional condition.	<b>show spantree nonforwardingreason</b> [port <i>port-string</i> ] [sid <i>sid</i> ]

## Example

The following example shows a switching device with Loop Protect enabled on port lag.0.2, SID 56:

```
Extremenetworks->show spantree lp port lag.0.2 sid 56
    LoopProtect is enabled on port lag.0.2, SID 56
Extremenetworks->show spantree lplock port lag.0.2 sid 56
    LoopProtect Lock status for port lag.0.2, SID 56_ is UNLOCKED
Extremenetworks->show spantree lpcapablepartner port lag.0.2
    Link partner of port lag.0.2_is LoopProtect-capable.
Extremenetworks->show spantree nonforwardingreason port lag.0.2
    Port lag.0.2 has been placed in listening or blocking state on SID 0 by the
    LoopProtect feature.
```

## Terms and Definitions

Table 18-9 lists terms and definitions used in Spanning Tree configuration.

**Table 18-9 Spanning Tree Terms and Definitions**

Term	Definition
Alternate port	Acts as an alternate path to the root bridge than that provided by the <a href="#">root port</a> .
Backup port	Acts as an backup for the path provided by a designated port toward the leaves of the Spanning Tree. Backup ports can exist only where two ports are connected together in a loopback mode or bridge with two or more connections to a shared LAN segment.
BID	Bridge identification, which is derived from the bridge's MAC address and bridge priority. The bridge with the lowest BID becomes the root bridge.
BPDU	Bridge Protocol Data Unit messages. Used by STP to exchange information, including designating a bridge for each switched LAN segment, and one root bridge for the Spanning Tree.
Bridge	Switching device.
Bridge priority	Assigns the bridge's relative priority compared to other bridges.
CIST	Common and Internal Spanning Tree created by MSTP to represent the connectivity of the entire network. This is equivalent to the single Spanning Tree used for STP and RSTP. Communications between MST regions occurs using the CIST.
CST	A Spanning Tree defined in the IEEE 802.1q standard that assumes one Spanning Tree instance for the entire bridged network, regardless of the number of VLANs.
Designated port	A forwarding port within an active topology elected for every switched LAN segment.
Edge port	Port on the edge of a bridged LAN.
FID	Filter Identifier. Each VLAN is associated to a FID. VLANs are mapped to SIDs using their FID association.
Forward delay	Time interval (in seconds) the bridge spends in listening or learning mode before it begins forwarding BPDUs.
Hello time	Time interval (in seconds) at which the bridge sends BPDUs.
ISL	Inter-Switch Link.



**Table 18-9 Spanning Tree Terms and Definitions (continued)**

Term	Definition
IST	A Spanning Tree instance that extends the CST inside the MST region and represents the entire MST region as a single CST virtual bridge to the outside world.
Loop Protect	Prevents or short circuits loop formation in a network with redundant paths by requiring ports to receive type 2 BPDUs (RSTP/MSTP) on point-to-point inter-switch links (ISLs) before their states are allowed to become forwarding.
Master port	The MSTI port whose connecting CIST port is root port for an entire MST region.
Max age	Maximum time (in seconds) the bridge can wait without receiving a configuration message (bridge “hello”) before attempting to reconfigure.
MST region	An MSTP group of devices configured together to form a logical region. The MST region presents itself to the rest of the network as a single device, which simplifies administration.
MSTI	Multiple Spanning Tree Instance. See <a href="#">Table 18-4</a> on page 22 for MSTI support per platform.
Path cost	Sum of the port costs in the best path to the root bridge.
Port cost	Value assigned to a port based on the speed of the port. The faster the speed, the lower the cost. This helps to determine the quickest path between the root bridge and a specified destination. The segment attached to the root bridge normally has a path cost of zero.
Port priority	Assigns a port’s priority in relation to the other ports on the same bridge.
Restricted TCN	Restricts the propagation of Topology Change Notices on a specified port when set to true.
Restricted Role	Disallows root role on a specified port when set to true.
Root bridge	Logical center of the Spanning Tree, used by STP to determine which paths to block and which to open.
Root port	Port in an active topology through which the root bridge can be reached.
SID	Spanning tree identifier. By default, SID 0 is assumed. VLANs are mapped to SIDs using their FID association.
SpanGuard	Prevents Spanning Tree respins that can occur when BPDUs are received on user ports and notifies network management that they were attempted.
TCN	Topology Change Notification.

## Shortest Path Bridging (SPB) Configuration

This chapter provides information about configuring and monitoring Shortest Path Bridging (SPB) on 7100-Series devices.

For information about...	Refer to page...
<a href="#">Using Shortest Path Bridging (SPB) in Your Network</a>	19-1
<a href="#">Implementing Shortest Path Bridging</a>	19-3
<a href="#">Shortest Path Bridging VLAN Configuration Overview</a>	19-4
<a href="#">Configuring Shortest Path Bridging VLAN</a>	19-6
<a href="#">Terms and Definitions</a>	19-7

### Using Shortest Path Bridging (SPB) in Your Network

Shortest Path Bridging (SPB), IEEE 802.1aq, is a protocol that provides data traffic a shortest cost path between any pair of switches in the SPB network. SPB features dynamic route calculation in a loop-free Layer-2 network and fast convergence time using IS-IS. The 7100-Series supports Shortest Path Bridging VLAN (SPBV).

SPB is administratively enabled by default, but requires that Spanning Tree be configured with the Spanning Tree version set to SPT, and the MST configuration name must be the same on all devices for a given SPB region. You can configure multiple SPB regions by assuring that each device within a given SPB region has the same MST configuration name. Spanning Tree defaults to version 3 (MSTP).

SPB uses IS-IS link state to populate the network topology information and calculate the Shortest Path Trees (SPTs). No IS-IS configuration is required. IS-IS is automatically enabled with a default NET as part of SPB region setup. The SPB capable bridge provides user access to an SPB region. A bridge joins an SPB region by forming an adjacency with a neighboring bridge through the exchange of IS-IS PDUs.

The base-VLAN is the customer VLAN that ingresses the SPBV network. It is used to manage operations in the SPT and provides access to the SPBV network. The Shortest Path VID (SPVID) is a VLAN used to identify a base-VLAN and SPT within the SPB network. Once a packet ingresses the SPB network, the SPVID is used. When the packet egresses the SPB network, the SPVID is translated back to the base-VLAN.


There are two types of ports in an SPBV region. Packets ingress and egress the SPB region on boundary ports. The boundary port is a member of the base-VLAN. Internal ports face into the SPB region and are members of the SPVID. An SPVID pool is administratively configured. The VLANs in the SPVID pool can only be used as SPVIDs. All SPB region ports must be administratively enabled for SPB.

Each base-VLAN at the boundary port is administratively assigned one of 16 Equal Cost Tree (ECT) algorithms as defined in IEEE 802.1Qaq. Each base-VLAN within the same SPB region must be assigned the same ECT algorithm. When traffic ingresses the SPB region on the base-VLAN, the packet VLAN (base-VLAN) is translated to an SPVID. By default this translation is dynamic, but you can optionally administratively set the base-VLAN to SPVID mapping.

Each VLAN on a bridge is associated with an SPVID whose tree is rooted at that bridge. This unique mapping is enforced by ISIS. As the traffic passes through the SPBV network, the SPVID is used. The base-VLAN and SPVID follow normal MAC and FID learning by making use of shared VLAN learning, but its FID will be defined by the mapped base-VLAN. Since the SPVID is unique to a device and base-VLAN, the SPVID defines the source bridge of traffic received within the SPBV region. The SPVID pool must be a size equal to or greater than the number of base-VLANs times the number of nodes in the SPB region that switch base-VLAN traffic. The ECT determines the next switch in the path through the SPB network. When the packet egresses the SPB network at a boundary port, the SPVID is translated back to the appropriate VLAN before being forwarded.

Figure 19-1 on page 19-2 shows an example where a user is trying to access server C in the Data Center through use of a spanning tree switched network. Spanning tree determines bridge A is the root of the CIST, and is blocking ports on bridges C, D, and E. The red line would be a normal MSTP switched traffic path, which makes an extra hop through router A. The shortest path is skipped due to the blocked port. The SPB network comprises a set of trees, each tree rooted at a switch where traffic ingresses the region. No ports are blocked and ISIS will create a shortest path through B and E to server C on the blue line.

**Figure 19-1 SPB Overview**




SPB provides benefits in addition to lower latency due to fewer hops. Use of IS-IS allows for quicker network recovery times than traditional Spanning Tree. SPB also provides for a more even distribution of traffic so switches, such as bridge A in Figure 19-1, do not become bottlenecks. SPB makes use of Equal Cost Trees (ECTs) to further refine traffic distribution.

Traffic can be mapped to access the SPBV backplane with different base-VLANs. A base-VLAN with one ECT algorithm may make a different path choice than another base-VLAN with a different ECT algorithm when faced with paths of equal cost. A network can be set up where traffic received on specific VLANs outside the SPBV network will all map to a single base-VLAN and will use the same ECT path, for example the blue line in [Figure 19-2](#) on page 19-3. Other VLANs can be mapped to a different base-VLAN and use a different ECT path, as indicated by the red line. All devices that are associated with a particular base-VLAN must be configured with the same ECT algorithm.

SPBV provides shortest paths without requiring changes in MTU size because it performs VID translation rather than adding its own header. It also fully supports congestion notification.

**Figure 19-2 SPBV Using Equal Cost Trees**



## Implementing Shortest Path Bridging

To implement SPBV:

1. Configure Spanning Tree on all devices in the region:
  - Set the Spanning Tree version to SPT on all devices in the SPB region
  - Configure the same MST configuration name on all devices in the SPB region
2. Configure an SPVID pool for this SPB region (same VLAN range for all devices in the region).
3. Assign the base-VLANs that will be used to ingress and egress the SPB region to SID **4093** or SID **spbv**.
4. Enable SPB on all ports that will take part in the SPB region.
5. Optionally, assign the desired ECT algorithm to each configured base-VLAN (unless the default algorithm is desired).
6. Optionally, administratively assign the base-VLAN to SPVID mapping for the base-VLAN on each device in the SPB region. When administratively assigning the base-VLAN to SPVID mapping, change the SPB VLAN mode to manual.

## Shortest Path Bridging VLAN Configuration Overview

For information about...	Refer to page...
<a href="#">SPBV Spanning Tree Configuration</a>	19-4
<a href="#">SPVID Pool</a>	19-4
<a href="#">Assigning a Base-VLAN to Use SPB</a>	19-5
<a href="#">Base-VLAN Configuration</a>	19-5
<a href="#">SPB Ports</a>	19-5

### SPBV Spanning Tree Configuration

Spanning Tree must be configured for SPBV to be operational. Configure Spanning Tree as documented in the *Extreme Networks 7100-Series Configuration Guide*. There are two important considerations when configuring Spanning Tree for SPBV:

- Spanning Tree version
- MST configuration name

Spanning Tree version defaults to MSPT. SPBV requires Spanning Tree version SPT. Use the **set spantree version** command in any configuration mode to set the Spanning Tree version to SPT. The following example sets the Spanning Tree version to SPT:

```
7100 Chassis(rw)->set spantree version spt
7100 Chassis(rw)->
```

The Spanning Tree MST configuration name must be the same for all devices in the SPBV region. The MST configuration name is set using the **set spantree mstcfgid** command in any configuration mode. The following example sets the MST configuration name to **spbv1** for this device:

```
7100 Chassis(rw)->set spantree mstcfgid cfgname spbv1
7100 Chassis(rw)->
```

### SPVID Pool

Shortest Path VLANs (SPVIDs) are the VLANs used by the SPBV region internal ports. An SPVID pool reserves a set of VLANs for SPVID use only. The number of VLANs reserved in the SPVID pool must be equal to or greater than the number of base-VLANs times the number of nodes in the SPBV region that switch base-VLAN traffic. For example if there are two base-VLANs, twenty nodes in the region, but only 16 forward base-VLAN-tagged traffic, the minimum number of VLANs needed in the SPVID pool is  $2 \times 16$  or 32 VLANs.

Use the **set spantree mstmap** command to specify a range of VLANs to reserve for the SPVID pool. The keyword **spvid** specifies SID 4095 for this Spanning Tree instance.



**Note:** The SPVID pool configuration must be the same for all nodes in a given SPBV region.

The following example sets the SPVID pool range to VLANs 2000 through 3000 for this device.

```
7100 Chassis(rw)->set spantree mstmap 2000-3000 sid spvid
7100 Chassis(rw)->
```

## Assigning a Base-VLAN to Use SPB

Base-VLANs within the Spanning Tree domain are assigned to use SPB by changing the SID to 4093. The keyword **spbv** can be used in place of 4093.

Use the **set spantree mstmap** command to specify one or more Base-VLANs to use SPB.

The following example configures Base-VLANs 100, 200, 300, and 400 to use SPB.

```
7100 Chassis(rw)->set spantree mstmap 100,200,300,400 sid spbv
7100 Chassis(rw)->
```

## Base-VLAN Configuration

There are two aspects to base-VLAN configuration:

- ECT algorithm assignment
- Base-VLAN to SPVID mapping

The base-VLAN belongs to the boundary port where packets ingress and egress the SPBV region. The base-VLAN manages operations in the SPT and provides access to the SPB network. A given base-VLAN in an SPBV region must be assigned the same ECT algorithm. The ECT determines the next switch in the path through the SPB. There are 16 ECT algorithms defined in IEEE 802.1Qaq.

Use the **set spb basevid ect-alg** command in any command mode to assign an ECT algorithm to a base-VLAN.

The following example assigns ECT algorithm 2 to base-VLAN 100:

```
7100 Chassis(rw)->set spb basevid 100 ect-alg ieee 2
7100 Chassis(rw)->
```

By default, base-VLANs are dynamically mapped to a unique SPVID from the SPVID pool for each node in the SPBV region. You can administratively map the base-VLAN to an SPVID for each node in the SPBV region. To administratively configure the base-VLAN to SPVID mapping:

- Set the SPB VLAN mode to manual using the **set spb system mode-vlan** command in any command mode
- Map the SPVID to the base-VLAN using the **spvid** option of the **set spb basevid ect-alg** command

The following example maps SPVID 2000 to base-VLAN 100:

```
7100 Chassis(rw)->set spb system mode-vlan manual
7100 Chassis(rw)->set spb basevid 100 ect-alg ieee 2 spvid 2000
7100 Chassis(rw)->
```

## SPB Ports

SPB must be enabled on all ports in the SPBV region. SPB defaults to disabled on ports. Use the **set spb port status enable** command in any command mode to enable SPB on all ports accessing the SPBV region.

The following example enables SPB on ports **ge.1.1** through **ge.1.5**:

```
7100 Chassis(rw)->set spb port ge.1.1-5 status enable
7100 Chassis(rw)->
```

## Configuring Shortest Path Bridging VLAN

This section provides a table of Shortest Path Bridging default values and a procedure for configuring a Shortest Path Bridging system.

Table 19-1 lists Shortest Path Bridging default values.

**Table 19-1 Default Shortest Path Bridging Parameters**

Parameter	Description	Default Value
SPB Status	Shortest Path Bridging device global state.	enabled
SPB Port Status	Shortest Path Bridging port state.	disabled
SPB System VLAN Mode	Configures the mode that determines whether the base-VID to SPVID mapping is manual or dynamic (auto).	auto
SPB Digest-Convention	Configures the SPB agreement digest convention.	loopfreeboth
Spanning Tree Version	Configures the Spanning Tree version; SPB requires SPT version 4	MSTP (version 3)

Procedure 19-1 describes Shortest Path Bridging configuration on the Extreme Networks 7100-Series devices. All commands used to configure Shortest Path Bridging can be entered in any command mode.

**Procedure 19-1 Configuring Shortest Path Bridging**

Step	Task	Command(s)
1.	Configure Spanning Tree on all SPB region devices.	See the <i>Extreme Networks 7100-Series Configuration Guide</i> Spanning Tree Configuration information.
2.	Enable SPB globally on the device. SPB is globally enabled by default.	<b>set spb status {enable   disable}</b>
3.	Set the Spanning Tree version to SPT on all devices in the SPB region.	<b>set spantree version spt</b>
4.	Configure the same Spanning Tree MST configuration name on all devices in the SPB region.	<b>set spantree mstcfgid {[cfgname name] [rev level]}</b>
5.	Configure an SPVID pool for this SPB region.	<b>set spantree mstmap spvid-list sid {4095   spvid}</b>
6.	Configure base-VLANs to use SPB.	<b>set spantree mstmap baseVid-list sid {4093   spbv}</b>
7.	Assign the ECT algorithm for each base-VLAN on this device.	<b>set spb basevid baseVid ect-alg ieee ect-alg</b>
8.	Optionally, set SPB system parameters: IS-IS area, digest convention, VLAN mode, or system ID.	<b>set spb system [area-address isis-area] [digest-convention {off   loopfreeboth}] [mode-vlan {auto   manual}]</b>
9.	Optionally, administratively assign the base-VLAN to SPVID mapping for the base-VLAN on each device in the SPB region.	<b>set spb basevid baseVid spvid spVid</b>
10.	Enable SPB on all ports that will take part in the SPB region.	<b>set spb port port-string status enable</b>

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

## Terms and Definitions

Table 19-2 lists terms and definitions used in this security mode configuration discussion.

**Table 19-2 Shortest Path Bridging Configuration Terms and Definitions**

Term	Definition
Shortest Path Bridging (SPB)	Shortest Path Bridging (SPB), defined in IEEE 802.1aq, is a protocol that provides data traffic a shortest cost path between any pair of switches in the SPB network, and features dynamic route calculation in a loop-free Layer-2 network, and fast convergence time using IS-IS.
Base-VLAN	The VLAN located on the SPBV boundary port that manages operations in the SPT and provides access to the SPBV network.
Shortest Path VID (SPVID)	The VLAN located on the port internal to the SPB region used to identify a base-VLAN and SPT within the SPB network.
SPB Boundary Port	The SPB port packets ingress and egress the SPB region on located on the base-VLAN.
SPB Internal Port	A port that faces into the SPB region and is a member of the SPVID.
Equal Cost Tree (ECT) algorithm	One of 16 algorithms defined in the IEEE 802.1Qaq standard used to select a path among available equal cost paths.





## VLAN Configuration

This chapter provides the following information about configuring and monitoring 802.1Q VLANs on Extreme Networks 7100-Series devices.

For information about...	Refer to page...
<a href="#">Using VLANs in Your Network</a>	20-1
<a href="#">Implementing VLANs</a>	20-2
<a href="#">Understanding How VLANs Operate</a>	20-3
<a href="#">VLAN Support on Extreme Networks 7100-Series Switches</a>	20-6
<a href="#">Configuring VLANs</a>	20-9
<a href="#">Terms and Definitions</a>	20-18
<a href="#">VLAN Provider Bridges</a>	20-19



**Note:** This document describes the configuration and operation of VLANs as defined by the IEEE 802.1Q standard and assumes that all devices being configured support that standard. No other types of VLANs will be covered.

### Using VLANs in Your Network

A VLAN is a Virtual Local Area Network — a grouping of network devices that is logically segmented by functions, project teams, or applications without regard to the physical location of users. For example, several end stations might be grouped as a department, such as Engineering or Finance, having the same attributes as a LAN, even though they are not all on the same physical LAN segment.


To accomplish this logical grouping, the network administrator uses 802.1Q VLAN-capable switching devices and assigns each switch port in a particular group to a VLAN. Ports in a VLAN share broadcast traffic and belong to the same broadcast domain. Broadcast traffic in one VLAN is not transmitted outside that VLAN.

Virtual LANs allow you to partition network traffic into logical groups and control the flow of that traffic through the network. Once the traffic and, in effect, the users creating the traffic, are assigned to a VLAN, then broadcast and multicast traffic is contained within the VLAN and users can be allowed or denied access to any of the network's resources. Also, you have the option of configuring some or all of the ports on a device to allow frames received with a particular VLAN ID and protocol to be transmitted on a limited number of ports. This keeps the traffic associated with a particular VLAN and protocol isolated from the other parts of the network.

The primary benefit of 802.1Q VLAN technology is that it allows you to localize and segregate traffic, improving your administrative efficiency, and enhancing your network security and performance.

Figure 20-1 shows a simple example of using port-based VLANs to achieve these benefits. In this example, two buildings house the Sales and Finance departments of a single company, and each building has its own internal network. The end stations in each building connect to a switch on the bottom floor. The two switches are connected to one another with a high speed link.

**Figure 20-1 VLAN Business Scenario**



Without any VLANs configured, the entire network in the example in Figure 20-1 would be a broadcast domain, and the switches would follow the IEEE 802.1D bridging specification to send data between stations. A broadcast or multicast transmission from a Sales workstation in Building One would propagate to all the switch ports on Switch A, cross the high speed link to Switch B, and then be propagated out all switch ports on Switch B. The switches treat each port as being equivalent to any other port, and have no understanding of the departmental memberships of each workstation.

Once Sales and Finance are placed on two separate VLANs, each switch understands that certain individual ports or frames are members of separate workgroups. In this environment, a broadcast or multicast data transmission from one of the Sales stations in Building One would reach Switch A, be sent to the ports connected to other local members of the Sales VLAN, cross the high speed link to Switch B, and then be sent to any other ports and workstations on Switch B that are members of the Sales VLAN. Separate VLANs also provides unicast separation between Sales and Finance. Finance can not ping Sales unless there is a routed VLAN configured for both Finance and Sales.

Another benefit to VLAN use in the preceding example would be your ability to leverage existing investments in time and equipment during company reorganization. If, for instance, the Finance users change location but remain in the same VLAN connected to the same switch port, their network addresses do not change, and switch and router configuration is left intact.

## Implementing VLANs

By default, all Extreme Networks switches run in 802.1Q VLAN operational mode. All ports on all Extreme Networks switches are assigned to a default VLAN (VLAN ID 1), which is enabled to operate and assigns all ports an egress status of untagged. This means that all ports will be allowed to transmit frames from the switch without a VLAN tag in their header. Also, there are no forbidden ports (prevented from transmitting frames) configured.

You can use the CLI commands described in this document to create additional VLANs, to customize VLANs to support your organizational requirements, and to monitor VLAN configuration.

## Preparing for VLAN Configuration

A little forethought and planning is essential to a successful VLAN implementation. Before attempting to configure a single device for VLAN operation, consider the following:

- What is the purpose of my VLAN design? (For example: security or traffic broadcast containment).
- How many VLANs will be required?
- What stations (end users, servers, etc.) will belong to them?
- What ports on the switch are connected to those stations?
- What ports will be configured as dynamic VLAN GARP VLAN Registration Protocol (GVRP) or Multiple VLAN Registration Protocol (MVRP) aware ports?
- What VLANs will not perform dynamic VLAN GVRP or MVRP processing?

Determining how you want information to flow and how your network resources can be best used to accomplish this will help you customize the tasks described in this document to suit your needs and infrastructure.

Once your planning is complete, you would proceed through the steps described in “[Configuring VLANs](#)” on page 20-9.

## Understanding How VLANs Operate

802.1Q VLAN operation differs slightly from how a switched networking system operates. These differences are due to the importance of keeping track of each frame and its VLAN association as it passes from switch to switch, or from port to port within a switch.

VLAN-enabled switches act on how frames are classified into a particular VLAN. Sometimes, VLAN classification is based on tags in the headers of data frames. These VLAN tags are added to data frames by the switch as the frames are transmitted out certain ports, and are later used to make forwarding decisions by the switch and other VLAN aware switches. In the absence of a VLAN tag header, the classification of a frame into a particular VLAN depends upon the configuration of the switch port that received the frame.

For information about...	Refer to page...
<a href="#">Learning Modes and Filtering Databases</a>	20-3
<a href="#">VLAN Assignment and Forwarding</a>	20-4
<a href="#">Example of a VLAN Switch in Operation</a>	20-6

## Learning Modes and Filtering Databases

Addressing information the switch learns about a VLAN is stored in the filtering database assigned to that VLAN. This database contains source addresses, their source ports, and VLAN IDs, and is referred to when a switch makes a decision as to where to forward a VLAN tagged frame. Each filtering database is assigned a Filtering Database ID (FID). The FID a VLAN belongs to can be displayed using the **show vlan** command.

A switch learns and uses VLAN addressing information by the following modes:

- **Independent Virtual Local Area Network (VLAN) Learning (IVL):** Each VLAN uses its own filtering database. Transparent source address learning performed as a result of incoming VLAN traffic is not made available to any other VLAN for forwarding purposes. This setting is useful for handling devices (such as servers) with NICs that share a common MAC address. One FID is assigned per VLAN. The FID value is the same as the VID it is assigned to. This is the default mode on Extreme Networks switches.
- **Shared Virtual Local Area Network (VLAN) Learning (SVL):** Two or more VLANs are grouped to share common source address information. This setting is useful for configuring more complex VLAN traffic patterns, without forcing the switch to flood the unicast traffic in each direction. This allows VLANs to share addressing information. It enables ports or switches in different VLANs to communicate with each other (when their individual ports are configured to allow this to occur). One FID is used by two or more VLANs. The FID value defaults to the lowest VID in the filtering database.

The VLAN learning mode for the switch and the assignment of multiple VLANs to a FID are configured in VLAN constraints using the **set vlan constraint** command. See “Appendix F” of the *IEEE Std 802.1Q™2011* standard for a detailed discussion of shared and independent VLAN learning modes.

## VLAN Assignment and Forwarding

### Receiving Frames from VLAN Ports

By default, Extreme Networks switches run in 802.1Q operational mode, which means that every frame received by the switch must belong to, or be assigned to, a VLAN. The type of frame under consideration and the filter setting of the switch determines how it forwards VLAN frames. This involves processing traffic as it enters (ingresses) and exits (egresses) the VLAN switch ports as described below.

#### Untagged Frames

When, for example, the switch receives a frame from Port 1 and determines the frame does not currently have a VLAN tag, but recognizes that Port 1 is a member of VLAN A, it will classify the frame to VLAN A. In this fashion, all untagged frames entering a VLAN switch assume membership in a VLAN.



**Note:** A VLAN ID is always assigned to a port. By default, it is the default VLAN (VLAN ID = 1).

The switch will now decide what to do with the frame, as described in “[Forwarding Decisions](#)” on page 20-5.

#### Tagged Frames

When, for example, the switch receives a tagged frame from Port 4 and determines the frame is tagged for VLAN C, it will classify it to that VLAN regardless of its port VLAN ID (PVID). This frame may have already been through a VLAN aware switch, or originated from a station capable of specifying a VLAN membership. If a switch receives a frame containing a tag, the switch will classify the frame in regard to its tag rather than the PVID for its port, following the ingress precedence rules listed below.

#### Ingress Precedence

VLAN assignment for received (ingress) frames is determined by the following precedence:

1. 802.1Q VLAN tag (tagged frames only)
2. Policy or Traffic Classification (which may overwrite the 802.1Q VLAN tag) For more information, refer to “[Configuring Protocol-Based VLAN Classification](#)” on page 20-15.
3. Port VID (PVID)

## Forwarding Decisions

VLAN forwarding decisions for transmitting frames is determined by whether or not the traffic being classified is or is not in the VLAN’s forwarding database as follows:

- **Unlearned traffic:** When a frame’s destination MAC address is not in the VLAN’s forwarding database (FDB), it will be forwarded out of every port on the VLAN’s egress list with the frame format that is specified. Refer to “[Broadcasts, Multicasts, and Unlearned Unicasts](#)” below for an example.
- **Learned traffic:** When a frame’s destination MAC address is in the VLAN’s forwarding database, it will be forwarded out of the learned port with the frame format that is specified. Refer to “[Learned Unicasts](#)” below for an example.

### Broadcasts, Multicasts, and Unlearned Unicasts

If a frame with a broadcast, multicast, or other unknown address is received by an 802.1Q VLAN aware switch, the switch checks the VLAN classification of the frame. The switch then forwards the frame out all ports that are identified in the Forwarding List for that VLAN. For example, if Port 3, shown in the example in [Figure 20-2](#), received the frame, the frame would then be sent to all ports that had VLAN C in their Port VLAN List.

### Learned Unicasts

When a VLAN switch receives a frame with a known MAC address as its destination address, the action taken by the switch to determine how the frame is transmitted depends on the VLAN, the VLAN associated FID, and if the port identified to send the frame is enabled to do so.

When a frame is received it is classified into a VLAN. The destination address is looked up in the FID associated with the VLAN. If a match is found, it is forwarded out the port identified in the lookup if, and only if, that port is allowed to transmit frames for that VLAN. If a match is not found, then the frame is flooded out all ports that are allowed to transmit frames belonging to that VLAN.

## Adding a MIB-II Interface Entry to a VLAN

A VTAP interface provides the data source input of a port mirror or SMON statistics collection. VTAP creation is the mechanism for adding a MIB-II interface table entry for a VLAN. When creating a VTAP interface, the specified VLAN is assigned a MIB-II ifIndex. A VLAN will not have a MIB-II ifIndex if a VTAP interface does not exist for it. Use the set vlan interface command to create a VTAP interface.

This example shows how to create a non-volatile MIB-II interface entry mapped to VLAN 1:

```
7100 Chassis(rw)->set vlan interface 1 create
7100 Chassis(rw)->show vlan interface 1
VLAN MIB-II Interfaces
```

```
Max Interfaces      : 16
Current Interfaces  : 1
```

```
VLAN      Port      Storage Type
```

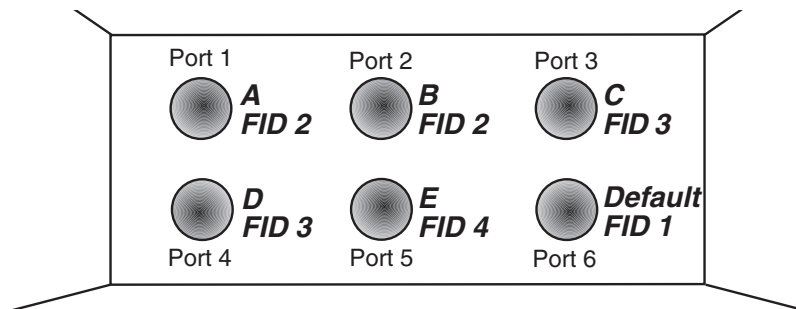
```
-----
1      vtap.0.1      non-volatile
7100 Chassis (rw) ->
```

## Example of a VLAN Switch in Operation

The operation of an 802.1Q VLAN switch is best understood from a point of view of the switch itself. To illustrate this concept, the examples that follow view the switch operations from *inside* the switch.

Figure 20-2 depicts the inside of a switch with six ports, numbered 1 through 6. The switch has been configured to associate VLAN A and B with FID 2, VLAN C and D with FID 3, and VLAN E with FID 4. It shows how a forwarding decision is made by comparing a frame's destination MAC to the FID to which it is classified.

**Figure 20-2 Inside the Switch**



Assume a unicast untagged frame is received on Port 3 in the example in Figure 20-2. The frame is classified for VLAN C (the frame's PVID is VLAN C). The switch would make its forwarding decision by comparing the destination MAC address to information previously learned and entered into its filtering database. In this case, the MAC address is looked up in the FDB for FID 3, which is associated with VLANs C and D. Let's say the switch recognizes the destination MAC of the frame as being located out Port 4.

Having made the forwarding decision based on entries in the FID, the switch now examines the port VLAN egress list of Port 4 to determine if it is allowed to transmit frames belonging to VLAN C. If so, the frame is transmitted out Port 4. If Port 4 has not been configured to transmit frames belonging to VLAN C, the frame is discarded.

If, on the other hand, a unicast untagged frame is received on Port 5, it would be classified for VLAN E. Port 5 has its own filtering database and is not aware of what addressing information has been learned by other VLANs. Port 5 looks up the destination MAC address in its FID. If it finds a match, it forwards the frame out the appropriate port, if and only if, that port is allowed to transmit frames for VLAN E. If a match is not found, the frame is flooded out all ports that are allowed to transmit VLAN E frames.

## VLAN Support on Extreme Networks 7100-Series Switches

For information about...	Refer to page...
<a href="#">Maximum Active VLANs</a>	20-7
<a href="#">Configurable Range</a>	20-7
<a href="#">VLAN Types</a>	20-7

For information about...	Refer to page...
<a href="#">Dynamic VLAN Support</a>	20-8

## Maximum Active VLANs

The total number of active VLANs supported on Extreme Networks 7100-Series switches is up to 4094.

## Configurable Range

The allowable user-configurable range for VLAN IDs (VIDs) on Extreme Networks 7100-Series switches is from 2 through 4094. This range is based on the following rules:

- **VID 0** is the null VLAN ID, indicating that the tag header in the frame contains priority information rather than a VLAN identifier. It cannot be configured as a port VLAN ID (PVID).
- **VID 1** is designated the default PVID value for classifying frames on ingress through a switched port. This default can be changed on a per-port basis.
- **VID 4095** is reserved by IEEE for implementation use.



**Notes:** Each VLAN ID in a network must be unique. If you enter a duplicate VLAN ID, the Extreme Networks switch assumes you intend to modify the existing VLAN.

## VLAN Types

Extreme Networks switches support traffic classification for the following VLAN types:

### Static and Dynamic VLANs

All VLANs on an Extreme Networks switch are categorized as being either static or dynamic. Static VLANs are those that are explicitly created on the switch itself, persistently remaining as part of the configuration, regardless of actual usage. Dynamic VLANs, on the other hand, are not necessarily persistent. Their presence relies on the implementation of GVRP or MVRP and its effect on egress membership as described in [“Dynamic VLAN Support”](#) on page 20-8.

### Port-Based VLANs

Port-based VLANs are configured by associating switch ports to VLANs in two ways: first, by manipulating the port VLAN ID (PVID); and second, by adding the port itself to the egress list of the VLAN corresponding to the PVID. Any traffic received by a port is associated to the VLAN identified by the port's PVID. By virtue of this association, this traffic may egress the switch only on those ports listed on the VLAN's egress list. For example, given a VLAN named “Marketing,” with an ID value of 6, by changing the PVID values of ports 1 through 3 to 6, and adding those ports to the egress list of the VLAN, we effectively restrict the broadcast domain of Marketing to those three ports. If a broadcast frame is received on port 1, it will be transmitted out ports 2 and 3 only. In this sense, VLAN membership is determined by the location of traffic ingress, and from the perspective of the access layer—where users are most commonly located—egress is generally untagged.

### Policy-Based VLANs

Rather than making VLAN membership decisions simply based on port configuration, each incoming frame can be examined by the classification engine which uses a match-based logic to



assign the frame to a desired VLAN. For example, you could set up a policy which designates all e-mail traffic between the management officers of a company to a specific VLAN so that this traffic is restricted to certain portions of the network. With respect to network usage, the administrative advantages of policy classification would be application provisioning, acceptable use policy, and distribution layer policy. All of these provisions may involve simultaneous utilization of inter-switch links by multiple VLANs, requiring particular attention to tagged, forbidden, and untagged egress settings.

As described above, PVID determines the VLAN to which all untagged frames received on associated ports will be classified. Policy classification to a VLAN takes precedence over PVID assignment if:

- policy classification is configured to a VLAN, and
- PVID override has been enabled for a policy profile, and assigned to port(s) associated with the PVID.

For more information, refer to the Policy Classification chapter.

## Dynamic VLAN Support

Dynamic VLAN support automatically creates VLANs across a switched network by dynamically establishing and updating a device's knowledge of the set of VLANs that currently have active members. Two dynamic VLAN capable protocols are supported on the 7100-Series:

- GARP (Generic Attribute Registration Protocol) VLAN Registration Protocol (GVRP)
- Multiple VLAN Registration Protocol (MVRP)

MVRP is a replacement protocol for GVRP based upon the Multiple Registration Protocol (MRP) defined by the IEEE 802.1ak amendment to the IEEE 802.1Q standard. MVRP uses a single packet event propagation for all 4094 VLANs, providing a more efficient encoding of PDUs which reduces the amount of traffic generated by the protocol. MVRP also performs filter database flushing on a per-VID basis during a topology change.

In a GVRP context, Spanning Tree performs filter database flushing on a per-port basis during a topology change. MVRP's ability to localize topology changes result in a more rapid healing of network failures without interrupting services to unaffected VLANs.


By default, both GVRP and MVRP are enabled both globally and on all VLANs, and disabled at the port level. For dynamic VLAN creation to occur, either MVRP or GVRP must be enabled globally as well as on each individual port. GVRP and MVRP can not both be enabled on the same port. Dynamic VLAN creation can be configured on both a per port and per VLAN basis.

### How Dynamic VLAN Support Works

When a VLAN has egress, the information is transmitted out ports on the device in a GVRP or MVRP formatted frame, using the GVRP or MVRP multicast MAC address. A switch that receives this frame examines the frame and extracts the VLAN IDs. The dynamic VLAN protocol then dynamically registers (creates) the VLANs and adds the receiving port to its tagged member list for the extracted VLAN IDs. The information is then transmitted out the other GVRP or MVRP configured ports of the device.

**Figure 20-3** shows an example of how VLAN Blue from end station A would be propagated across a switch network. In this figure, port 1 of Switch 4 is registered as being a member of VLAN Blue and Switch 4 declares this fact out all its ports (2 and 3) to Switch 1 and Switch 2. These two switches register this in the port egress lists of the ports (Switch 1, port 1 and Switch 2, port 1) that received the frames with the information. Switch 2, which is connected to Switch 3 and Switch 5 declares the same information to those two switches and the port egress list of each port is updated with the new information, accordingly.

Figure 20-3 Example of VLAN Propagation Using GVRP or MVRP



**Note:** If a port is set to “forbidden” for the egress list of a VLAN, then the VLAN’s egress list will not be dynamically updated with that port.

Administratively configuring a VLAN on an 802.1Q switch creates a static VLAN entry that will always remain registered and will not time out. However, GVRP or MVRP dynamically created entries will time out, and their registrations will be removed from the member list if the end station is removed. This ensures that, if switches are disconnected or if end stations are removed, the registered information remains accurate.

The end result of the dynamic VLAN configuration is that each port’s egress list is updated with information about VLANs that reside on that port, even if the actual station on the VLAN is several hops away.

## Configuring VLANs

Once you have planned your implementation strategy as described in “[Preparing for VLAN Configuration](#)” on page 20-3, you can begin configuring VLANs as described in this section.

For information about...	Refer to page...
<a href="#">Default Settings</a>	20-10
<a href="#">Configuring Static VLANs</a>	20-11
<a href="#">Creating a Secure Management VLAN</a>	20-13
<a href="#">Configuring Dynamic VLANs</a>	20-14
<a href="#">Configuring Protocol-Based VLAN Classification</a>	20-15
<a href="#">Configuring IGMP VLAN Snooping</a>	20-17
<a href="#">Monitoring VLANs</a>	20-18

## Default Settings

Table 20-1 lists VLAN parameters and their default values.

**Table 20-1 Default VLAN Parameters**

Parameter	Description	Default Value
garp or mrp timer	Configures the three GARP or MRP timers. These timers are critical and should only be modified by someone familiar with the 802.1Q standard.	<ul style="list-style-type: none"> <li>Join timer: 20 centiseconds</li> <li>Leave timer: 60 centiseconds</li> <li>Leaveall timer: 1000 centiseconds</li> </ul>
gvrp	Enables or disables the GARP VLAN Registration Protocol (GVRP) on a specific set of ports or all ports for dynamic VLAN creation. GVRP must be enabled when using GVRP for dynamic VLAN creation. GVRP can not be enabled on ports MVRP is enabled on.	<ul style="list-style-type: none"> <li>Disabled at the port level</li> <li>Enabled at the global level</li> </ul>
mvrp	Enables or disables the Multiple VLAN Registration Protocol (MVRP) on a specific set of ports or all ports. MVRP must be enabled when using MVRP for dynamic VLAN creation. MVRP can not be enabled on ports GVRP is enabled on.	<ul style="list-style-type: none"> <li>Disabled at the port level</li> <li>Enabled at the global level</li> </ul>
vlan restricted	Disables dynamic VLAN creation on a VLAN basis when VLAN restricted is enabled on the VLAN.	Dynamic VLAN is enabled on all VLANs. VLAN restricted feature is disabled.
IGMP last member query interval	Configures the last member query interval. This is the maximum response time inserted into group-specific queries which are sent in response to Leave Group messages. It is also the amount of time between group-specific query messages.	1 second
IGMP VLAN max response time	Configures the maximum query response time (in tenths of a second).	100 deciseconds (10 seconds)
IGMP VLAN query interval	Configures the frequency (in seconds) of host-query frame transmissions.	125 seconds
IGMP VLAN robustness	Configures the robustness value.	2
IGMP VLAN version	Selects the IGMP version. Options are version 1 or version 2.	Version 2
port discard	Ports can be set to discard frames based on whether or not they contain a VLAN tag.	No frames are discarded
port ingress filter	When enabled on a port, the VLAN IDs of incoming frames are compared to the port's egress list. If the received VLAN ID does not match a VLAN ID on the port's egress list, the frame is dropped.	Disabled

**Table 20-1 Default VLAN Parameters (continued)**

Parameter	Description	Default Value
port vlan ID (PVID)	802.1Q VLAN/port association.	VLAN1/ Default VLAN
vlan constraint	Configures VLANs to use an independent or shared filtering database.	VLANs use an independent filtering database
vlan dynamicegress	Enables or disables dynamic egress processing for a given VLAN.	Disabled
vlan egress	Configures the egress ports for a VLAN and the type of egress for the ports. Egress type can be tagged, untagged, or forbidden.	Tagged
vlan name	Associates a text name to one or more VLANs.	None

## Configuring Static VLANs

[Procedure 20-1](#) describes how to create and configure a static VLAN. Unspecified parameters use their default values.

### Procedure 20-1 Static VLAN Configuration

Step	Task	Command(s)
1.	Show existing VLANs.	<b>show vlan</b>
2.	Create VLAN. Valid values are <b>1–4094</b> . Each <i>vlan-id</i> must be unique. If an existing <i>vlan-id</i> is entered, the existing VLAN is modified.	<b>set vlan create <i>vlan-id</i></b>
3.	Optionally, assign a name to the VLAN. Valid strings are from 1 to 32 characters.	<b>set vlan name <i>vlan-id string</i></b>
4.	Assign switched ports to the VLAN. This sets the port VLAN ID (PVID). The PVID determines the VLAN to which all untagged frames received on the port will be classified.	<b>set port vlan <i>port-string vlan-id</i></b>




**Note:** If the VLAN specified has not already been created, the above command will create it. It will also add the VLAN to the port's egress list as untagged, and remove the default VLAN from the port's egress list. This automatically changes the existing untagged VLAN egress permission to match the new PVID value.

**Procedure 20-1 Static VLAN Configuration (continued)**

Step	Task	Command(s)
5.	<p>Configure VLAN egress, which determines which ports a frame belonging to the VLAN may be forwarded out on.</p> <p><b>Static configuration:</b> Add the port to the VLAN egress list for the device.</p> <ul style="list-style-type: none"> <li>The default setting, <b>tagged</b>, allows the port to transmit frames for a particular VLAN.</li> <li>The <b>untagged</b> setting allows the port to transmit frames without a VLAN tag. This setting is usually used to configure a port connected to an end user device.</li> <li>The <b>forbidden</b> setting prevents the port from participating in the specified VLAN and ensures that any dynamic requests for the port to join the VLAN will be ignored.</li> </ul> <p>If necessary, remove ports from the VLAN egress list.</p> <ul style="list-style-type: none"> <li>If specified, the <b>forbidden</b> setting will be cleared from the designated ports and the ports will be reset as allowed to egress frames, if so configured by either static or dynamic means.</li> <li>If <b>forbidden</b> is not specified, tagged and untagged egress settings will be cleared from the designated ports.</li> </ul> <p><b>Dynamic configuration:</b> By default, dynamic egress is disabled on all VLANs. If dynamic egress is enabled for a VLAN, the device will add the port receiving a frame to the VLAN's egress list as untagged according to the VLAN ID of the received frame.</p>	<p><b>set vlan egress</b> <i>vlan-id port-string forbidden   tagged   untagged</i></p> <p><b>clear vlan egress</b> <i>vlan-list port-string [forbidden]</i></p> <p><b>set vlan dynamic</b><i>egress</i> <i>vlan-id {enable   disable}</i></p>
6.	<p>Optionally, set VLAN constraints to control the filtering database a VLAN will use for forwarding traffic. Filtering databases can be shared or independent. By default, filtering databases are independent.</p>	<p><b>set vlan constraint</b> <i>vlan-id set-num [shared   independent]</i></p>
7.	<p>Optionally, enable ingress filtering on a port to drop those incoming frames that do not have a VLAN ID that matches a VLAN ID on the port's egress list.</p>	<p><b>set port ingress-filter</b> <i>port-string enable</i></p>
8.	<p>Optionally, choose to discard tagged or untagged, (or both) frames on selected ports. Select <b>none</b> to allow all frames to pass through.</p>	<p><b>set port discard</b> <i>port-string {tagged   untagged   none   both}</i></p>
9.	<p>If the device supports routing, enter interface configuration mode and configure an IP address on the VLAN interface.</p>	<p><b>configure</b> <b>interface</b> <i>vlan vlan-id</i> <b>ip address</b> <i>ip-address ip-mask</i> <b>no shutdown</b></p>

**Procedure 20-1 Static VLAN Configuration (continued)**

Step	Task	Command(s)
	<b>Note:</b> Each VLAN interface must be configured for routing separately using the interface command shown above. To end configuration on one interface before configuring another, type <b>exit</b> at the command prompt. Enabling interface configuration mode is required for completing interface-specific configuration tasks.	

**Example Configuration**

The following shows an example 7100-Series device configuration using the steps in [Procedure 20-1](#). In this example, VLAN 100 is created and named VLANRED. Ports ge.1.2, 1.3 and 1.4 are assigned to VLAN 100 and added to its egress list. VLAN 100 is then configured as a routing interface with an IP address of 120.20.20.24.

```
7100 Chassis(rw)->set vlan create 100
7100 Chassis(rw)->set vlan name 100 VLANRED
7100 Chassis(rw)->set port vlan ge.1.2-4 100
    The PVID is used to classify untagged frames as they
    ingress into a given port. Would you like to add the selected
    port(s) to this VLAN's untagged egress list and remove them
    from all other VLANs untagged egress list (y/n) [n]? y
    NOTE: Choosing 'y' will not remove the port(s) from previously
    configured tagged egress lists.
7100 Chassis(rw)->router
7100 Chassis(rw)-router->configure terminal
7100 Chassis(rw)-router-config->interface vlan 100
7100 Chassis(rw)-router-config-intf-Vlan-100->ip address 120.20.20.1/24
7100 Chassis(rw)-router(config-intf-Vlan-100)->no shutdown
```

If you want to configure a port to drop incoming frames that do not have a VLAN ID that matches a VLAN ID on the port's egress list, use the **set port ingress-filter** command. For example:

```
7100 Chassis(rw)->set port ingress-filter ge.1.2-4 enable
```

If you want to configure a port to discard tagged or untagged incoming frames, use the **set port discard** command. For example, to configure the ports to drop tagged frames on ingress:

```
7100 Chassis(rw)->set port discard ge.1.2-4 tagged
```

**Creating a Secure Management VLAN**

If you are configuring an Extreme Networks device for multiple VLANs, it may be desirable to configure a management-only VLAN. This allows a station connected to the management VLAN to manage the device. It also makes management secure by preventing configuration through ports assigned to other VLANs.

[Procedure 20-2](#) provides an example of how to create a secure management VLAN. This example, which sets the new VLAN as VLAN 2, assumes the management station is attached to ge.1.1, and wants untagged frames. The process described in this section would be repeated on every switch device that is connected in the network to ensure that each switch has a secure management VLAN.

**Procedure 20-2 Secure Management VLAN Configuration**

Step	Task	Command(s)
1.	Create a new VLAN.	<b>set vlan create 2</b>
2.	Set the PVID for the host port and the desired switch port to the VLAN created in Step 2.	<b>set port vlan host.0.1; ge.1.1 2</b>
3.	If not done automatically when executing the previous command, add the host port and desired switch port(s) to the new VLAN's egress list.	<b>set vlan egress 2 host.0.1; ge.1.1 2 untagged</b>
4.	Set a private community name to assign to this VLAN for which you can configure access rights and policies.	<b>set snmp community private</b>



**Note:** By default, community name—which determines remote access for SNMP management—is set to **public** with read-write access. For more information, refer to your device's SNMP documentation.

## Configuring Dynamic VLANs

[Procedure 20-3](#) describes the configuration of dynamic VLANs using the GARP (Generic Attribute Registration Protocol) VLAN Registration Protocol (GVRP). [Procedure 20-4](#) on page 20-15 describes the configuration of dynamic VLANs using the Multiple VLAN Registration Protocol (MVRP). Both GVRP and MVRP are enabled globally and disabled at the port level. A port can only be enabled for GVRP or MVRP, but not both. Whichever dynamic VLAN protocol you are using, it must be globally enabled and also enabled on specific ports in order to generate and process the appropriate GVRP or MVRP advertisement frames. The dynamic VLAN protocol must be enabled on the VLAN for the appropriate advertisement frames to be processed. Dynamic VLAN creation can be disabled per VLAN using the VLAN restriction feature supported by both GVRP and MVRP.



**Note:** Refer to “[Dynamic VLAN Support](#)” on page 20-8 for conceptual information about both GVRP and MVRP.



**Caution:** The setting of GARP or MRP timers is critical and should only be changed by personnel familiar with 802.1Q standards.

**Procedure 20-3 GVRP Configuration**

Step	Task	Command(s)
1.	Show existing GVRP configuration for a port or list of ports. If no <i>port-string</i> is entered, the global GVRP configuration and all port GVRP configurations are displayed.	<b>show gvrp [port-string]</b>
2.	Enable GVRP on those ports assigned to a VLAN. GVRP is disabled at the port level by default.	<b>set gvrp enable port-string</b>
3.	Optionally, disable GVRP processing on specified VLANs.	<b>set gvrp vlan vlan-list restricted disable</b>
4.	Display GVRP status for system VLANs.	<b>show gvrp vlan {vlan-list   all} restricted</b>
5.	Display the existing GARP timer values.	<b>show garp timer [port-string]</b>

**Procedure 20-3 GVRP Configuration (continued)**

Step	Task	Command(s)
6.	Optionally, set the GARP join, leave, and leaveall timer values. Each timer value is in centiseconds.	<b>set garp timer</b> {[join timer-value] [leave timer-value] [leaveall timer-value]} port-string

**Procedure 20-4 MVRP Configuration**

Step	Task	Command(s)
1.	Show existing MVRP configuration for a port or list of ports. If no <i>port-string</i> is entered, the global MVRP configuration and all port MVRP configurations are displayed.	<b>show mvrp</b> [port-string]
2.	Enable MVRP on those ports assigned to a VLAN. MVRP is disabled at the port level by default.	<b>set mvrp enable</b> port-string
3.	Optionally, disable MVRP processing on specified VLANs.	<b>set mvrp vlan</b> vlan-list <b>restricted disable</b>
4.	Display MVRP status for system VLANs.	<b>show mvrp vlan</b> {vlan-list   all} <b>restricted</b>
5.	Display the existing MRP timer values.	<b>show mrp timer</b> [port-string]
6.	Optionally, set the MRP join, leave, and leaveall timer values or enable or disable the periodic timer which has a fixed interval of 1 second. Each timer value is in centiseconds.	<b>set mrp timer</b> {[join timer-value] [leave timer-value] [leaveall timer-value] [periodic {enable   disable}] port-string

## Configuring Protocol-Based VLAN Classification

Protocol-based VLANs can be configured using the policy classification CLI commands, as shown in this section, or NetSight Policy Manager.

[Procedure 20-5](#) describes how to define protocol-based frame filtering policies to assign frames to particular VLANs. Refer to your Extreme Networks policy configuration and CLI documentation for more information.



**Note:** Depending on your Extreme Networks switching device, your options for configuring policy classification may differ from the examples provided in this section. Refer to your device's documentation for a list of CLI commands and functions supported.

**Procedure 20-5 Configuring Protocol-Based VLAN Classification**

Step	Task	Command(s)
1.	Create the VLANs to which frames will be assigned by the policy. Valid values are <b>1–4094</b> .	<b>set vlan create</b> vlan-id



**Procedure 20-5 Configuring Protocol-Based VLAN Classification (continued)**

Step	Task	Command(s)
2.	Configure VLAN egress, which determines which ports a frame belonging to the VLAN may be forwarded out on. The default setting, <b>tagged</b> , allows the port to transmit frames for a particular VLAN. The <b>untagged</b> setting allows the port to transmit frames without a VLAN tag. The <b>forbidden</b> setting disallows egress on the port.	<b>set vlan egress</b> <i>vlan-id port-string</i> [ <b>forbidden</b>   <b>tagged</b>   <b>untagged</b> ]
3.	Disable ingress filtering on the ingress ports on which the policy will be applied. Disabled is the default ingress filtering setting.	<b>set port ingress-filter</b> <i>port-string</i> <b>disable</b>
4.	Create the policy profile that enables PVID override. This function allows a policy rule classifying a frame to a VLAN to override PVID assignment configured with the <b>set port vlan</b> command. When none of its associated classification rules match, the configuration of the policy profile itself will determine how frames are handled by default. In this case, the default VLAN is specified with the <b>pvid</b> <i>pvid</i> parameter.	<b>set policy profile</b> <i>profile-index</i> [ <b>name</b> <i>name</i> ] [ <b>pvid-status</b> { <b>enable</b>   <b>disable</b> }] [ <b>pvid</b> <i>pvid</i> ]
5.	Configure the administrative rules that will assign the policy profile to all frames received on the desired ingress ports.	<b>set policy rule admin-profile</b> <b>port</b> <i>port-string</i> [ <b>port-string</b> <i>port-string</i> ] [ <b>admin-pid</b> <i>admin-pid</i> ]
6.	Configure the classification rules that will define the protocol to filter on and the VLAN ID to which matching frames will be assigned.	<b>set policy rule</b> <i>profile-index</i> { <i>protocol data</i> } [ <b>mask</b> <i>mask</i> ] [ <b>vlan</b> <i>vlan</i> ]

**Example Configuration**

The following shows an example 7100-Series device configuration using the steps in [Procedure 20-5](#). This example configures a policy that ensures that IP traffic received on the specified ingress ports will be mapped to VLAN 2, while all other types of traffic will be mapped to VLAN 3.

- Two VLANs are created: VLAN 2 and VLAN 3.
- Ports 1 through 5 on the Gigabit Ethernet IOM in slot 4 are configured as egress ports for the VLANs while ports 8 through 10 on the Gigabit Ethernet IOM in slot 5 are configured as ingress ports that will do the policy classification.
- Policy profile number 1 is created that enables PVID override and defines the default behavior (classify to VLAN 3) if none of the classification rules created for the profile are matched.
- Administrative rules are created that apply policy profile number 1 to all frames received on the ingress ports ge.5.8 through 10.
- Classification rules are created for policy profile number 1 that assign IP frames to VLAN 2. The rules identify IP frames by using the **ether** protocol parameter, which classifies on the Type field in the headers of Layer 2 Ethernet II frames, and the protocol data of 0x0800 (IP type), 0x0806 (ARP type), and 0x8035 (RARP type).

```

7100 Chassis(rw)->set vlan create 2,3
7100 Chassis(rw)->set vlan egress 2 ge.4.1-2
7100 Chassis(rw)->set vlan egress 3 ge.4.3-5
7100 Chassis(rw)->set port ingress-filter ge.5.8-10 disable
7100 Chassis(rw)->set policy profile 1 name protocol_based_vlan pvid-status enable
    pvid 3
7100 Chassis(rw)->set policy rule admin-profile port ge.5.8 port-string ge.5.8
    admin-pid 1
7100 Chassis(rw)->set policy rule admin-profile port ge.5.9 port-string ge.5.9
    admin-pid 1
7100 Chassis(rw)->set policy rule admin-profile port ge.5.10 port-string ge.5.10
    admin-pid 1
7100 Chassis(rw)->set policy rule 1 ether 0x0800 mask 16 vlan 2
7100 Chassis(rw)->set policy rule 1 ether 0x0806 mask 16 vlan 2
7100 Chassis(rw)->set policy rule 1 ether 0x8035 mask 16 vlan 2

```

## Configuring IGMP VLAN Snooping

IGMP Layer 2 snooping allows the Extreme Networks switch for a specific VLAN to actively participate in IGMP traffic forwarding. IGMP snooping depends on the presence of an upstream IGMP querier. Whenever it receives an IGMP query, the switch forwards the query out the appropriate VLAN ports. IGMP snooping allows per-port traffic patterns in VLANs with multiple ports. It is disabled by default.

For more information, refer to the *Extreme Networks 7100-Series CLI Reference*.

[Procedure 20-6](#) describes how to configure IGMP snooping for a VLAN.

### Procedure 20-6 IGMP Snooping for a VLAN Configuration

Step	Task	Command(s)
1.	Enable IGMP snooping for a VLAN or a range of VLANs.	<b>set igmp enable</b> <i>vlan-id</i>
2.	Enable querying on this VLAN, and specify the IGMP querier source address.	<b>set igmp query-enable</b> <i>vlan-id</i> <b>address</b> <i>ip-address</i>
3.	Set the version of IGMP to use. Enter <b>1</b> for IGMPv1, <b>2</b> for IGMPv2, or <b>3</b> for IGMPv3.	<b>set igmp config</b> <i>vlan-id</i> <b>igmp-version</b> <b>1 2 3</b>
4.	Set the Last Member interval value, which can be 1–255.	<b>set igmp config</b> <i>vlan-id</i> <b>last-member-interval</b> <i>value</i>
5.	Set the Max Response Time which can be 1–255 seconds.	<b>set igmp config</b> <i>vlan-id</i> <b>max-response-time</b> <i>seconds</i>
6.	Set the Query Interval, which can be 1–65535 seconds.	<b>set igmp config</b> <i>vlan-id</i> <b>query-interval</b> <i>seconds</i>
7.	Set the Robustness value, which can be 2–255.	<b>set igmp config</b> <i>vlan-id</i> <b>robustness</b> <i>value</i>
8.	Optionally, create a static IGMP entry, or add ports to an existing entry. The entry can be in the form of an IP multicast address or IP group address.	<b>set igmp add-static</b> <i>{IP-multicast-address   IP-group-address vlan-id}</i> [ <b>modify</b> ] <i>port-string</i>

## Monitoring VLANs

Table 20-2 describes the **show** commands that display information about VLAN configurations. Refer to *Extreme Networks 7100-Series CLI Reference* for a description of the output of each **show** command.

**Table 20-2 Displaying VLAN Information**

Task	Command
Display all existing VLANs.	<b>show vlan</b>
Display the VLAN constraint setting.	<b>show vlan constraint</b> [vlan id]
Display the VLAN dynamic egress setting.	<b>show vlan dynamicegress</b> [vlan id]
Display all static VLANs.	<b>show vlan static</b>
Display ports assigned to VLANs.	<b>show port vlan</b> [port-string]
Display existing GVRP settings.	<b>show gvrp</b> [port-string]
Display existing MVRP settings.	<b>show mvrp</b> [port-string]
Display MVRP counters information.	<b>show mvrp counters</b> [port-string]
Display existing GVRP VLAN restricted settings.	<b>show gvrp vlan</b> {vlan-list   all} <b>restricted</b>
Display existing MVRP VLAN restricted settings.	<b>show mvrp vlan</b> {vlan-list   all} <b>restricted</b>
Display GARP timer values for one or more ports.	<b>show garp timer</b> [port-string]
Display MRP timer values for one or more ports.	<b>show mrp timer</b> [port-string]
Display IGMP VLAN configuration.	<b>show igmp config</b> [vlan id]
Display IGMP enable state of VLAN.	<b>show igmp enable</b> [vlan id]
Display all groups on a given VLAN.	<b>show igmp groups</b> [vlan id]
Display IGMP VLAN query state.	<b>show igmp query</b> [vlan id]
Display static ports on the given vid, group.	<b>show igmp static</b> [vlan id]

## Terms and Definitions

Table 20-3 lists terms and definitions used in VLAN configuration.

**Table 20-3 VLAN Terms and Definitions**

Term	Definition
Default VLAN	The VLAN to which all ports are assigned upon initialization. The default VLAN has a VLAN ID of 1 and cannot be deleted or renamed.
Filtering Database	A database structure within the switch that keeps track of the associations between MAC addresses, VLANs, and interface (port) numbers. The Filtering Database is referred to when a switch makes a forwarding decision on a frame.

**Table 20-3 VLAN Terms and Definitions (continued)**

Term	Definition
Filtering Database Identifier (FID)	Addressing information that the device learns about a VLAN is stored in the filtering database assigned to that VLAN. Several VLANs can be assigned to the same FID to allow those VLANs to share addressing information. This enables the devices in the different VLANs to communicate with each other when the individual ports have been configured to allow communication to occur.  The configuration is accomplished using the Local Management VLAN Forwarding Configuration screen. By default a VLAN is assigned to the FID that matches its VLAN ID.
Forwarding List	A list of the ports on a particular device that are eligible to transmit frames for a selected VLAN.
GARP Multicast Registration Protocol (GMRP)	A GARP application that functions in a similar fashion as GVRP, except that GMRP registers multicast addresses on ports to control the flooding of multicast frames.
GARP VLAN Registration Protocol (GVRP)	A GARP application used to dynamically create VLANs across a switched network.
Generic Attribute Registration Protocol (GARP)	GARP is a protocol used to propagate state information throughout a switched network.
Multiple VLAN Registration Protocol (MVRP)	An MRP application used to dynamically create VLANs across a switched network.
Multiple Registration Protocol (MRP)	MRP is a protocol used to propagate state information throughout a switched network.
Port VLAN List	A per port list of all eligible VLANs whose frames can be forwarded out one specific port and the frame format (tagged or untagged) of transmissions for that port. The Port VLAN List specifies what VLANs are associated with a single port for frame transmission purposes.
Tag Header (VLAN Tag)	Four bytes of data inserted in a frame that identifies the VLAN/frame classification. The Tag Header is inserted into the frame directly after the Source MAC address field. Twelve bits of the Tag Header represent the VLAN ID. The remaining bits are other control information.
Tagged Frame	A data frame that contains a Tag Header. A VLAN aware device can add the Tag Header to any frame it transmits.
Untagged Frame	A data frame that does not have a Tag Header.
VLAN ID	A unique number (between 1 and 4094) that identifies a particular VLAN.
VLAN Name	A 32-character alphanumeric name associated with a VLAN ID. The VLAN Name is intended to make user-defined VLANs easier to identify and remember.

## VLAN Provider Bridges

By extending VLAN tagging technology via the IEEE 802.1Q-2011 standard, an organization managing a service provider network can provide external user groups (departments, customers) to each have their own separate logical network (a LAN consisting of multiple VLANs) through the provider network to predefined egress ports. The provider can assign a single VLAN through their network for all traffic egressing through a port, instead of having to create and manage a separate VLAN for each customer VLAN. This feature, commonly called Provider Bridges (also

VLAN stacking, or Q-in-Q), performs Layer 2 tunneling from one customer network location through the provider network to another customer network location.

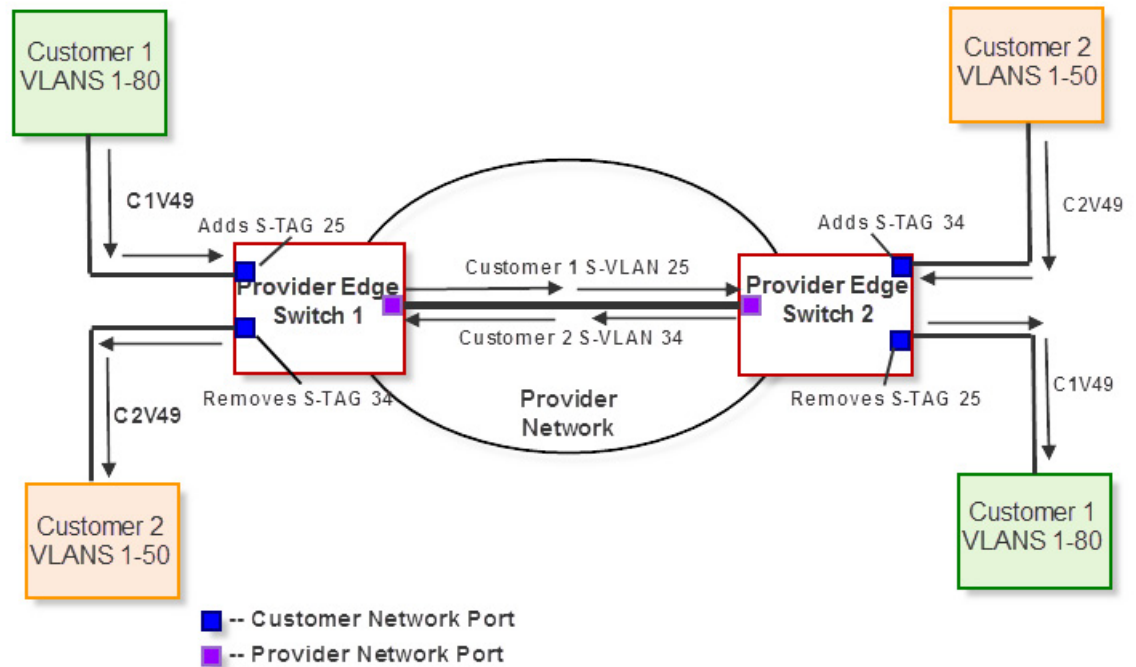
Provider bridges allow a service provider to assign customer traffic to a service instance called an S-VLAN (Service VLAN). When a packet from a customer VLAN arrives on an 7100-Series switch at the edge of the provider network, it is already identified by a tag called a C-TAG (customer tag). The edge switch encapsulates the packet with another tag called an S-TAG (service provider tag) for relay through the service provider network. The packet egresses from the provider network to the destination customer site through an egress port at which the S-TAG is stripped off. The packet continues on to the remote customer network guided by its original C-VLAN tagging, unaltered.

**Table 20-4 Provider Bridge VLAN Tags**

TAG Name	Description
C-TAG	Customer VLAN tag. 32 bit tag where the first 16 bits represent the ether type for the customer (0x8100) and the lower 16 bits represent the Priority Code Point (PCP), Canonical Frame Indicator (CFI), and customer VLAN ID (C-VID) associated with the packet.
S-TAG	Service provider VLAN tag. 32 bit tag where the first 16 bits represent the ether type for the service provider (0x88a8) and the lower 16 bits represent the PCP, DE (Drop Eligible), and service provider VLAN ID (S-VID) associated with the packet.

[Figure 20-4](#) illustrates the Provider Bridges function in a provider network with two customers, each with a campus on either side of the provider network and the need to connect their VLANs inexpensively, transparently, and securely. VLAN 49 from Customer 1 ingresses the provider network through the Customer Network Port (CNP) on Provider Edge Switch 1, which adds S-TAG 25 to packets with the C-VID for Customer 1 VLAN 49 (C1V49). After traversing the provider network as S-VLAN 25, the packets egress the provider network at the designated CNP on Provider Edge Switch 2, where the S-TAG 25 is removed and the packets are forwarded to their destination. Even though Customer 2 has a VLAN 49 (C2V49), this traffic is tagged with S-TAG 34 and traverses the provider network as S-VLAN 34.

Figure 20-4 Provider Bridges in Provider Network



If other traffic from Customer 1 (for example, VLAN 22) must traverse the provider network to get to the remote Customer 1 campus, the provider edge switches would add the S-TAG for S-VLAN 25 on C1V22 packets as well. In this way all Customer 1 traffic (80 VLANs) can pass over the provider network using a single S-VLAN, and arrive at their destinations with C-VLAN intact.

Table 20-5 defines the types of ports used in the Provider Bridge architecture. In hardware terms they are identical, but they differ in their roles in the bridging feature architecture.

Table 20-5 Provider Bridge-related Port Types

Port Type	Description
Customer Bridge Port	Customer bridge ports are the existing, default type of ports carrying C-VLAN traffic on provider edge switches. They interpret the C-TAG and relay packets using the C-VID.
Customer Network Port (CNPs)	<p>CNPs are ports resident on provider edge switches that connect the customer network to the provider network. Traffic received on a CNP is assigned to an S-VID based on the ingress settings of the CNP and encapsulated with an S-TAG. CNPs remove the S-TAG from the packets egressing the provider network. In cases where two provider networks are connected via CNPs, packets egress the device with the S-TAG. In these cases the ingress S-VID is translated to a relay S-VID which maps the S-VLANs of one provider network to the S-VLANs of the second provider network.</p> <p>Customer Network Ports are configured as untagged egress ports via the <b>set vlan egress</b> command.</p>
Provider Network Port (PNPs)	<p>PNPs interconnect switches within the provider network and operate using S-TAGs and S-VIDs for relay operation.</p> <p>Provider Network Ports are configured as tagged egress ports via the <b>set vlan egress</b> command.</p>

## Configuring Provider Bridges

There are two possible bridge modes for bridging customer traffic:

- **Customer Bridge Mode**

The default mode of routing external C-VLANs via C-TAGs.

- **Provider Bridge Mode**

Provider Bridge mode enables tunneling in provider network by adding S-TAGs to C-VLANs and transporting them as dedicated service VLANs.

### Customer Bridge Mode

The default mode for bridging and routing of external (customer) VLANs through an Extreme Networks 7100-Series switch is called Customer Bridging. In a customer bridge, all packets are bridged and routed based on the VLAN identified in the C-TAG. Routing through and egress from the network must be configured on a per-VLAN basis, so there must be a VLAN configured for each VLAN passing through the switch.

### Provider Bridge Mode

When the bridge mode is set to provider-bridge, switches in a service provider network relay packets based on service VLANs (using S-TAGs) rather than customer VLANs (using C-TAGs). The VLAN configuration now operates over the S-VLAN for relay operations instead of the C-VLAN. Customer-provided STP and MVRP addresses are preserved in the packets, but are ignored in the provider network. The provider network can be configured to use its own Spanning Tree and MVRP addresses to switch packets around the S-VLAN.

Any customer traffic received on a provider network ingress port will be placed in the S-VLAN defined for that port, and will be delivered to the customer network connected to the provider network egress port. For bidirectional VLAN traffic between the two remote customer networks, both of these provider ports must be configured on the same service VLAN.

[Procedure 20-7](#) shows an example configuration for transporting any number of received customer VLANs across a provider network in provider bridge mode. This configuration would apply to one of the switches in the provider network. A similar configuration must be executed on each provider switch in the path to the destination customer network. There must be an S-VLAN configured for each egress port to a customer network.

#### Procedure 20-7 Configuring a Provider Bridge

Step	Task	Command(s)
1.	Set the bridge mode to provider bridge.	<b>set bridge mode provider-bridge</b>
2.	Create an S-VLAN in the provider network. In this example, the new provider S-VLAN is 400.	<b>set vlan create 400</b>
3.	Add port ge.1.1 and ge.1.2 (both on this switch) to the S-VLAN 400 egress port list. This action sets the ports' VLAN ID to 400.	<b>set port vlan ge.1.1-2 400 modify-egress</b>
4.	Specify that port ge.1.1 can transmit frames out to the customer network without (VLAN 400) provider tagging.	<b>set vlan egress 400 ge.1.1 untagged</b>
5.	At port ge.1.2, the frames for VLAN 400 are egressed to another switch in the provider network <b>tagged</b> with a VLAN 400 S-TAG.	<b>set vlan egress 400 ge.1.2 tagged</b>

[Figure 20-5](#) illustrates the following example configuration:

**Provider Edge Switch 1:**

```

set bridge mode provider-bridge
set vlan create 400
set port vlan ge.1.1-2 400 modify-egress
set vlan egress 400 ge.1.1 untagged
set vlan egress 400 ge.1.2 tagged
set vlan create 500
set port vlan ge.1.2-3 500 modify-egress
set vlan egress 500 ge.1.3 untagged
set vlan egress 500 ge.1.2 tagged

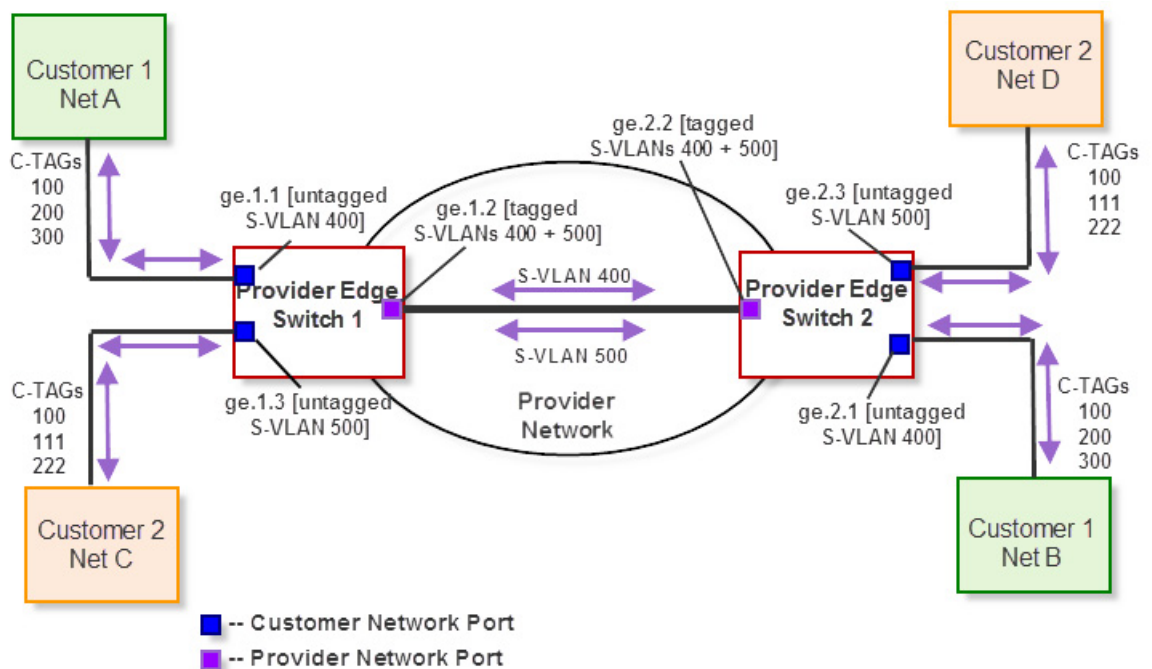
```

**Provider Edge Switch 2:**

```

set bridge mode provider-bridge
set vlan create 400
set port vlan ge.2.2-3 400 modify-egress
set vlan egress 400 ge.2.1 untagged
set vlan egress 400 ge.2.2 tagged
set vlan create 500
set port vlan ge.2.2-3 500 modify-egress
set vlan egress 500 ge.2.3 untagged
set vlan egress 500 ge.2.2 tagged

```

**Figure 20-5 Provider Bridge Configuration Example**

To clear provider bridge mode, revert your provider network by setting each bridge back to customer bridge mode. Use the **show bridge mode** command to view the current bridge mode.





## Link Aggregation Control Protocol (LACP) Configuration

This document describes the link aggregation feature and its configuration on Extreme Networks 7100-Series devices.

For information about...	Refer to page...
<a href="#">Using Link Aggregation in Your Network</a>	21-1
<a href="#">Implementing Link Aggregation</a>	21-2
<a href="#">Link Aggregation Overview</a>	21-3
<a href="#">Configuring Link Aggregation</a>	21-9
<a href="#">Link Aggregation Configuration Examples</a>	21-11
<a href="#">Terms and Definitions</a>	21-19

### Using Link Aggregation in Your Network

IEEE 802.3ad link aggregation provides a standardized means of grouping multiple parallel Ethernet interfaces into a single logical Layer 2 link. The formed group of Ethernet interfaces is referred to as a Link Aggregation Group (LAG). Dynamic LAG formation and activation is provided by the Link Aggregation Control Protocol (LACP).

Each pair of LAG physical ports is made up of a local port on the device responsible for LACP negotiation, referred to as the actor, and its directly linked remote port on the device participating in the LACP negotiation, referred to as the partner. LAGs form automatically based upon a set of criteria (see “[How a LAG Forms](#)” on page 21-3).

Only LAG members in the attached state carry user traffic. Once the LAG is formed, the system ID, made up of a system priority and the device MAC address, determines which device will be in charge of choosing the LAG port members that will be moved to the attached state. While port speed is not a criteria for joining a LAG, the port speed must match for all ports that are placed in the LACP attached state. Aggregatable ports not selected to carry traffic for this LAG are available to the next LAG as long as LAG resources are not depleted. Should LAG resources become depleted, aggregatable ports are placed in LACP standby state.

802.3ad LACP aggregations can be run between combinations of switches, routers, and edge devices, such as a server, that support LACP.



**Note:** Earlier (proprietary) implementations of port aggregation referred to groups of aggregated ports as “trunks”.

The concept of grouping multiple ports into a single link is not a new idea. Cabletron's SmartTrunk, Cisco's Inter Switch Link trunking, and Adaptec's Duralink are previous examples. The problem with these older methods, from the network administrators point of view, is that they are proprietary. Administrators who wanted to implement faster logical links faced major problems if they also wanted, or needed, to use a different brand of networking hardware. Link aggregation is standards based allowing for interoperability between multiple vendors in the network.

Older implementations required manual configuration. With LACP, if a set of links can aggregate, they will aggregate. LACP's ability to automatically aggregate links represents a timesaver for the network administrator who will not be required to manually configure the aggregates. However, manual overrides are provided for when the administrator needs to customize. Link aggregation also provides for rapid configuration and reconfiguration when there are changes in the physical connections. Link aggregation will automatically and quickly converge the new configuration. This convergence typically occurs in one second or less.

Link aggregation is a cost effective way to implement increased bandwidth. A major benefit of link aggregation is the ability to incrementally add bandwidth in a linear fashion. Without link aggregation, if there is a need to increase the bandwidth for a 100Mbps pipe, the only choice is an exponential upgrade to a 1000Mbps pipe. If there is a need for a 300Mbps pipe, aggregating three 100Mbps ports is both less expensive, because a forklift hardware upgrade is avoided, and makes for more efficient use of the system ports that are already available.

The physical links within the aggregate can serve as redundant backups to one another. Since only a single MAC address representing the entire aggregate is presented to the MAC client, the failure of any link within the aggregate is transparent. Failover is handled within the link aggregation sublayer.

## Implementing Link Aggregation

To implement link aggregation:

- Enable LACP on the network device
- Optionally set a non-default system priority for the device
- Optionally change the administratively assigned key for each port on the device
- Optionally enable single port LAGs on the device
- Enable LACP port active state on all ports that will take part in link aggregation
- Optionally change LAG parameters on each port
- Optionally change how flows will behave when changes take place to the LAG
- Optionally change the load balancing behavior for flows over the LAG
- Optionally assign static ports to a LAG when the partner device only supports a non-LACP method of aggregation

# Link Aggregation Overview

This section provides an overview of link aggregation configuration.

## LACP Operation

In order to allow LACP to determine whether a set of links connect to the same device, and to determine whether those links are compatible from the point of view of aggregation, it is necessary to be able to establish:

- A globally unique identifier for each device that participates in link aggregation.
- A means of identifying the set of capabilities associated with each port and with each aggregator, as understood by a given device.
- A means of identifying a LAG and its associated aggregator.

For each aggregatable port in the device, LACP:

- Maintains configuration information (reflecting the inherent properties of the individual links as well as those established by network administration) to control aggregation.
- Exchanges configuration information with other devices to allocate the link to a LAG.



**Note:** A given link is allocated to, at most, one LAG at a time. The allocation mechanism attempts to maximize aggregation, subject to management controls.

- Attaches the port to the aggregator used by the LAG, and detaches the port from the aggregator when it is no longer used by the LAG.
- Uses information from the partner device's link aggregation control entity to decide whether to aggregate ports.

The operation of LACP involves the following activities:

- Checking that candidate links can actually be aggregated.
- Controlling the addition of a link to a LAG and the creation of the group if necessary.
- Monitoring the status of aggregated links to ensure that the aggregation is still valid.
- Removing a link from a LAG if its membership is no longer valid, and removing the group if it no longer has any member links.

## How a LAG Forms

LAGs form automatically with LACP enabled on the device. There are four criteria for forming a LAG. Both actor and partner ports must:

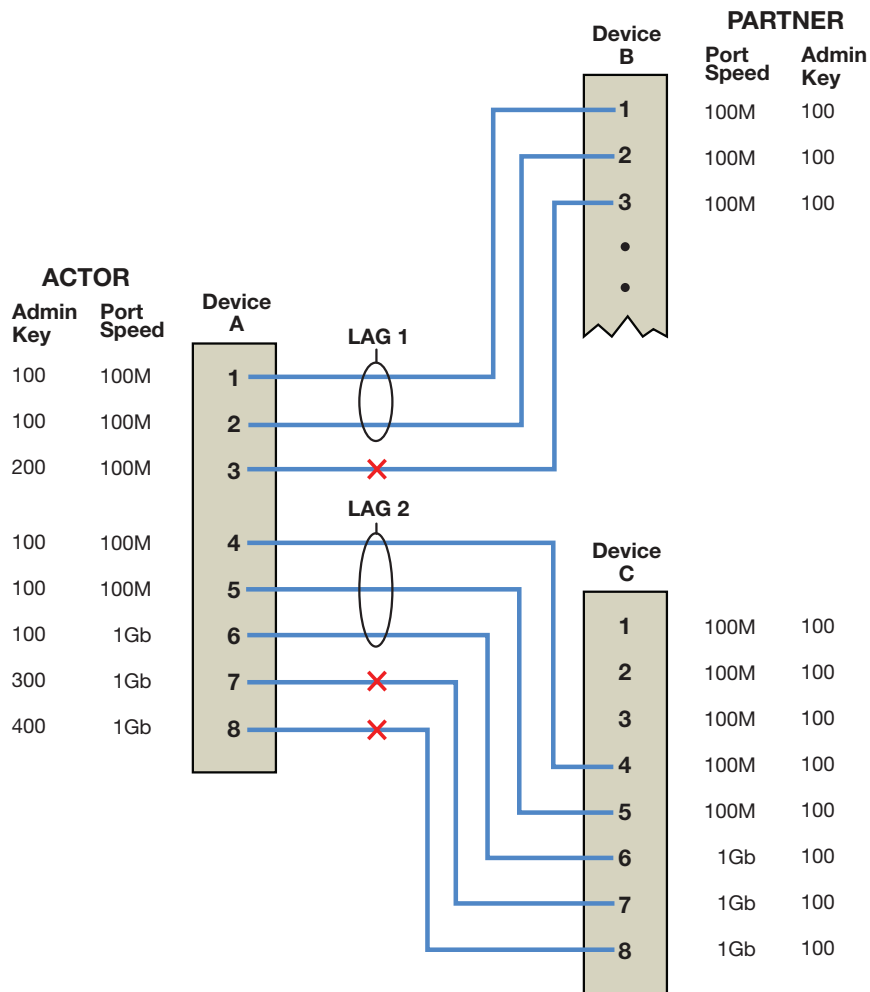
1. Operate in full duplex mode.
2. Have matching local LAG and physical port admin keys for the device controlling LAG formation.
3. Operate in parallel in that a LAG can have only two devices associated with it.
4. Consist of two or more physical actor to partner port pairings unless the single port LAG feature is enabled.

Figure 21-1 displays a LAG formation example containing three devices with five 100Mbps ports and three 1Gb ports configured. For this example, all ports are operating in full-duplex mode, and

the admin key for all LAG ports has been set to 100. Device A is the actor and therefore determines which ports will join a LAG. Devices B and C are the partners.

In our example two LAGs have formed because the actor ports are shared between two partner devices. Attempting to form a single LAG using all the actor ports would have broken the rule that actor and partner ports must operate in parallel.

**Figure 21-1 LAG Formation**



Actor ports 1 - 3 on device A directly connect to partner ports 1 - 3 on device B:

- We have already stated that all ports are operating in full-duplex mode, so rule 1 is satisfied for all three ports.
- Investigating the port admin keys, we see that ports 1 and 2 on device A are set to 100 (the same setting as all LAG ports on the device), while port 3 on device A is set to 200. Because the port admin keys are the same for both the LAG port and these physical ports, ports 1 and 2 satisfy rule 2. Because the admin key for physical port 3 is different from any possible LAG for this device, port 3 can not be part of any LAG.
- Because ports 1 and 2 for both the actor and partner operate in parallel with each other, rule 3 is satisfied for these ports.
- Rule 4 is satisfied, regardless of whether single port LAGs are enabled, because there are two aggregatable port pairings between devices A and B.

For these reasons, LAG 1 (lag.0.1) is formed using actor and partner ports 1 and 2.

Actor ports 4 - 8 on device A directly connect to partner ports 4 - 8 on device C:

- Because all ports are operating in full-duplex mode, rule one is satisfied for all five ports.
- Investigating port admin keys, we see that ports 4 - 6 on device A are set to 100 (the same setting as all LAG ports on the device), while ports 7 and 8 on device A are set to 300 and 400, respectively. Because port admin keys for all LAGs and the physical ports 4 - 6 are the same, physical ports 4 - 6 satisfy rule 2. Because the admin key settings for physical ports 7 and 8 do not agree with any LAG admin key setting on the device, ports 7 and 8 can not be part of any LAG.
- Because ports 4 - 6 for both the actor and partner operate in parallel with each other, rule 3 is satisfied for these ports.
- Rule 4 is satisfied, regardless of whether single port LAG is enabled, because there are three aggregatable port pairings between devices A and C.

For these reasons, LAG 2 is formed using actor and partner ports 4 - 6.



**Note:** Port speed is not a consideration in the forming phase for LAGs. LAG 2 contains 100Mbps and 1Gb port members.

## Attached Ports

Once a LAG is formed, two steps must take place before traffic can pass over the LAG:

- The device that will choose which ports to move to the attached state must be identified
- The process of moving the chosen ports to the LACP attached state must take place

A system ID, made up of the device MAC address and the system priority, is associated with each device. The device with the lower system priority is in charge of selecting the LAG members to move to the attached state. If a system priority tie occurs, the system with the lower MAC address value breaks the tie.

Only LAG members with the same port speed can be moved to the attached state. In a case where multiple speeds are present in a LAG, the LAG member with the lowest port priority on the device in charge, as well as all other members with the same port speed as the member with the lowest port priority, are selected and moved to the attached state. Using LAG2 in [Figure 21-1](#) on page 21-4 as an example, if the LAG2 member port priorities are set as shown in [Table 21-1](#) on page 21-5, ports 4 and 5 are moved to the attached state.

**Table 21-1 LAG2 Port Priority Assignments**

Port Number	Port Speed	Port Priority
4	100Mbps	200
5	100Mbps	300
6	1Gb	300

This is true because port 4 has the lowest priority of the three ports currently in the LAG, and port 5 has the same speed as the port with the lowest priority in the LAG, regardless of its priority. Because port 6 has both a different speed and a higher priority than the port with the lowest priority in the LAG, it is not moved to the attached state.

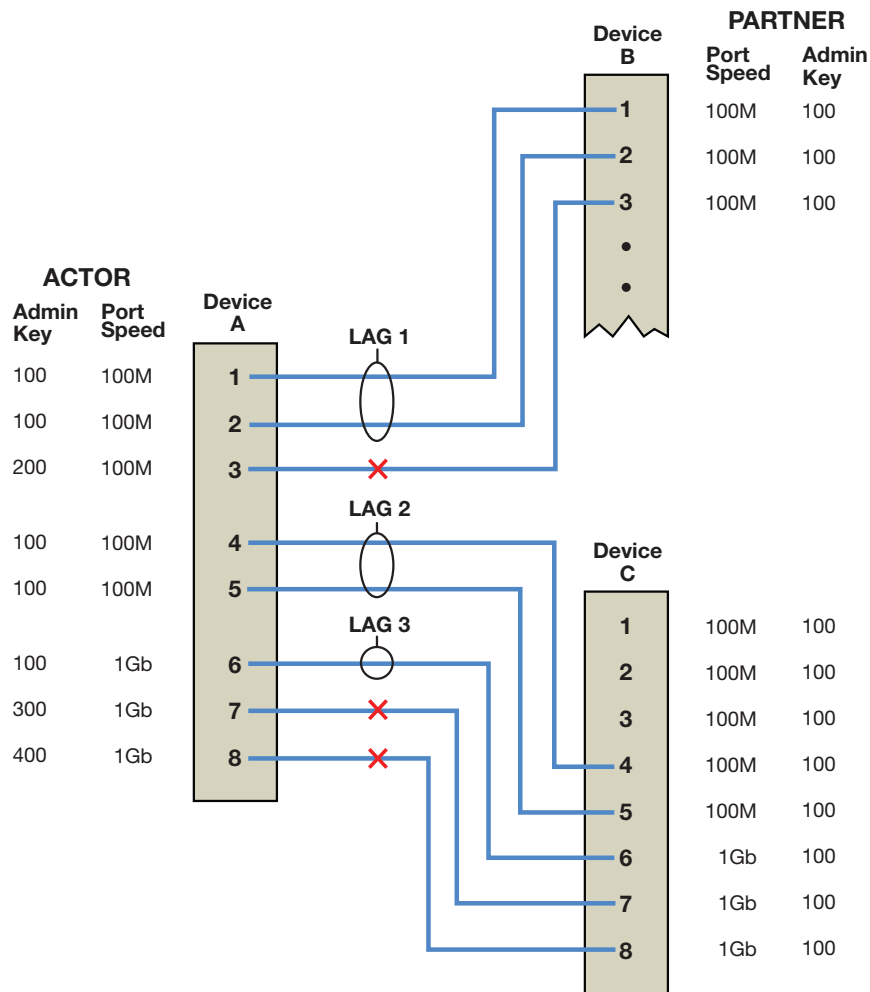
If LAG members with different port speeds should tie for the lowest port priority, the LAG member with the lowest port number breaks the tie. In our example, should all three ports have

the same port priority, ports 4 and 5 would still be the ports moved to the attached state because port 4 has the lowest port number and port 5 has the same port speed as port 4.

If in our example you wanted the reverse outcome of port 6 moved to the attached state instead of ports 4 and 5, setting port 6 to a lower priority than ports 4 and 5, as well as enabling the single port LAG feature on this device, would accomplish that goal.

Aggregatable ports not moved to the attached state are made available to form another LAG providing a LAG resource is available for this system. Port 6 in [Figure 21-1](#) on page 21-4, was not moved to the attached state. The only criteria port 6 does not meet to form its own LAG is rule 4: being a single aggregatable port. The single port LAG feature must be enabled for port 6 to form a LAG. If single port LAG is enabled on this system, port 6 would form and attach to LAG 3. [Figure 21-2](#) illustrates the three LAGs described in this example.

**Figure 21-2 LAGs Moved to Attached State**



Should an aggregatable port be available with all LAG resources depleted for this system, the port is placed in LACP standby state. Ports in standby state do not forward traffic. If all ports initially moved to the attach state for a given LAG become unavailable, a LAG resource will then be available. LACP will initiate a new selection process using the ports in standby state, using the same rules as the initial process of forming LAGs and moving ports to the attached state.

## Single Port Attached State Rules

By default, a LAG must contain two or more actor and partner port pairs for the LAG to be initiated by this device. A feature exists to allow the creation of a single port LAG that is disabled by default. If single port LAG is enabled, a single port LAG can be created on this device. If single port LAG is disabled, a single port LAG will not be initiated by this device. If a peer device is able to form a single port LAG and advertises its willingness to do so, a single port LAG can form.

There are three conditions under which a single port LAG can exist and the LAG member can be moved to the attached state:

- The single port LAG feature is enabled.
- or,
- The single port LAG feature is disabled, but the peer device is able and willing to form a single port LAG.
- or,
- An already existing LAG configuration persists through a device or module reset. If upon reset there is only a single port active for an already existing LAG, that single port will move to the attached state regardless of the single port LAG setting.

## LAG Port Parameters

LAG port parameters can be changed per port.

[Table 21-2](#) specifies the LACP port parameters that can be changed.

**Table 21-2 LAG Port Parameters**

Term	Definition
Port Admin Key	The port admin key can be set for both the actor and partner side of the link. The admin key only affects the local device. LACP uses this value to determine which underlying physical ports are capable of aggregating. Aggregator ports allow only underlying ports with physical port and LAG admin keys that match to join a LAG. Setting the physical port admin key to a different value than any LAG resource on the device will ensure that this link does not join a LAG. Valid values are <b>1 - 65535</b> . Default value is <b>32768</b> .
Port Priority	Port priority can be set for both the actor and partner side of the link. The port priority plays a role in determining which set of ports will move to the attached state and pass traffic. The lower port priority, for the port on the system in charge of selecting ports to move to the attached state, determines which ports will actually move to the attached state. If a LAG is made up of ports with different speeds, setting a lower port priority to ports with the desired speed for the LAG will ensure that those ports move to the attached state. Port priority is also used to determine which ports join a LAG if the number of ports available exceeds the number of ports supported for that device. Valid values are <b>0 - 65535</b> , with lower values designating higher priority. Default value is <b>32768</b> .



**Table 21-2 LAG Port Parameters (continued)**

Term	Definition
Administrative State	<p>A number of port level administrative states can be set for both the actor and partner ports. The following port administrative states are set by default:</p> <ul style="list-style-type: none"> <li>• <b>lacpactive</b> - Transmitting LACP PDUs is enabled.</li> <li>• <b>lacptimeout</b> - Transmitting LACP PDUs every 30 seconds. If this state is disabled, LACP PDUs are transmitted every 1 second. Note that the actor and partner LACP timeout values must agree.</li> <li>• <b>lacppagg</b> - Aggregation on this port is enabled.</li> <li>• <b>lacpsync</b> - Transition to synchronization state is allowed.</li> <li>• <b>lacpcollect</b> - Transition to collection state is allowed.</li> <li>• <b>lacpdist</b> - Transition to distribution state is allowed.</li> <li>• <b>lacpdef</b> - Transition to defaulted state is allowed.</li> <li>• <b>lacpexpire</b> - Transition to expired state is allowed.</li> </ul> <p><b>Notes:</b> It is recommended that you do not change these default states unless you know what you are doing. Contact Extreme Networks customer support should you need assistance modifying port level administrative states.</p>
Partner Default System ID	A default partner system ID can be set. This is a default MAC address for the system partner.
LACP PDU processing	(Optional) LACP PDU processing can be enabled or disabled for this port.

## The Out-Port Algorithm

The out-port algorithm determines the criteria to be used for data forwarding port selection. There are two algorithm criteria to choose from:

- Destination IP address and Source IP address (dip-sip). This is the most finely tuned criteria in that a port will be assigned based upon a specific IP address combination for the flow. All flows for this IP address combination transit the assigned physical port.
- Destination MAC address and Source MAC address (da-sa). This criteria is less finely tuned in that a port will be assigned based upon the MAC address combination for the flow. All flows for this MAC address combination transit the assigned port.

## Static Port Assignment

Static port assignment allows you to assign ports to a LAG when the partner device does not support LACP, but does support another proprietary form of link aggregation. To assign a static port, specify the LAG port ID, the admin key value for this LAG, and the ports to be assigned. If you do not specify an admin key value, a key will be assigned according to the specified aggregator. For example, a key of 4 would be assigned to lag.0.4.

## Platform LAG and Physical Port Support

The number of LAGs and the number of ports per LAG supported are platform specific. The number of LAGs supported is on a system basis. See [Table 21-3](#) for a listing of the number of LAGs and the number of ports per LAG supported for your platform.

## Configuring Link Aggregation

This section provides details for the configuration of link aggregation on the 7100-Series products. [Table 21-4](#) lists link aggregation parameters and their default values.

**Table 21-4 Default Link Aggregation Parameters**

Parameter	Description	Default Value
LACP State	Current state of LACP on the device.	Enabled
System Priority	LACP system priority for this device.	32768
Port Key	The Port Administrative Key (also referred to as operational key).	32768
Port Priority	Determines which ports move to the attached state when ports of different speeds form a LAG. Also determines which ports join a LAG if the ports available exceed the number of ports supported by the device.	32768
Single Port State	Allows or disallows a LAG to be created with a single port.	Disabled (disallows creation of a single port LAG)
LACP Port Active State	Port state providing for transmission of LACP PDUs.	Disabled
LACP Port Timeout State	Port state determining the frequency of LACP PDU transmission and period before declaring the partner LACP port down if no response is received.	30 second: frequency of LACP PDU transmission 90 seconds: period before declaring the partner port down

[Procedure 21-1](#) describes how to configure link aggregation.

**Procedure 21-1 Configuring Link Aggregation**

Step	Task	Command(s)
1.	In switch command mode, enable LACP on the device.	<b>set lacp {disable   enable}</b>
2.	Optionally, change the system priority for the device.	<b>set lacp asyspri value</b>
3.	Optionally, change the administratively assigned key for each aggregation on the device.	<b>set lacp adminkey port-string value</b>
4.	Optionally, enable single port LAGs on the device.	<b>set lacp singleportlag {enable   disable}</b>

**Procedure 21-1 Configuring Link Aggregation (continued)**

Step	Task	Command(s)
5.	Optionally, enable port active state for all LAG participating ports and modify the LAG port parameters. See <a href="#">Table 21-2</a> on page 21-7 for a description of port parameters.	<pre> <b>set port lacp port</b> <i>port-string</i> {   [aadminkey <i>aadminkey</i>] [aportpri <i>aportpri</i>]   [padminsyspri <i>padminsyspri</i>] [padminsysid <i>padminsysid</i>] [padminkey <i>padminkey</i>]   [padminportpri <i>padminportpri</i>] [padminport <i>padminport</i>]    [aadminstate {lacpactive   lacptimeout   lacpagg   lacpsync   lacpcollect   lacpdist   lacpdef   lacpexpire}]    [padminstate {lacpactive   lacptimeout   lacpagg   lacpsync   lacpcollect   lacpdist   lacpdef   lacpexpire}]    [enable   [disable]] } </pre>
6.	Optionally, change the out-port behavior for flows over the LAG.	<b>set lacp outportAlgorithm</b> { <i>dip-sip</i>   <i>da-sa</i> }
7.	Optionally, assign static ports to a LAG when the partner device only supports a non-LACP method of aggregation.	<b>set lacp static</b> <i>lagportstring</i> [ <i>key</i> ] <i>port-string</i>

[Table 21-5](#) describes how to manage link aggregation.

**Table 21-5 Managing Link Aggregation**

Task	Command
Reset LACP to the default state of enabled.	<b>clear lacp state</b>
Reset LACP system priority or admin key settings to the default values.	<b>clear lacp</b> {[ <i>asyspri</i> ] [ <i>aadminkey port-string</i> ]}
Remove specific static ports from an aggregation.	<b>clear lacp static</b> <i>lagportstring port-string</i>
Reset the single port LAG feature to the default value of disabled.	<b>clear lacp singleportlag</b>
Reset a link aggregation port setting to the default value for one or more ports. See <a href="#">Table 21-2</a> on page 21-7 for a description of port parameters.	<pre> <b>clear port lacp port</b> <i>port-string</i> {   [aadminkey] [aportpri] [padminsyspri] [padminsysid]   [padminkey] [padminportpri] [padminport]    [aadminstate {lacpactive   lacptimeout   lacpagg   lacpsync   lacpcollect   lacpdist   lacpdef   lacpexpire   all}]    [padminstate {lacpactive   lacptimeout   lacpagg   lacpsync   lacpcollect   lacpdist   lacpdef   lacpexpire   all}] } </pre>
Reset the LACP out-put algorithm setting to its default value of DIS-SIP.	<b>clear lacp outportAlgorithm</b>

[Table 21-6](#) describes how to display link aggregation information and statistics.

**Table 21-6 Displaying Link Aggregation Information and Statistics**

Task	Command
Display the global LACP enable state, or display information about one or more aggregator ports.	<b>show lacp</b> [ <b>state</b>   <i>port-string</i> ]
Display the status of the single port LAG function.	<b>show lacp singleportlag</b>
Display link aggregation information for one or more underlying physical ports.	<b>show port lacp port</b> <i>port-string</i> <b>{[status {detail   summary}]   [counters]}</b> <b>[sort {port   lag}]</b>
Display the current configured out-port algorithm.	<b>show lacp outportAlgorithm</b>

## Link Aggregation Configuration Examples

This section presents two configuration examples:

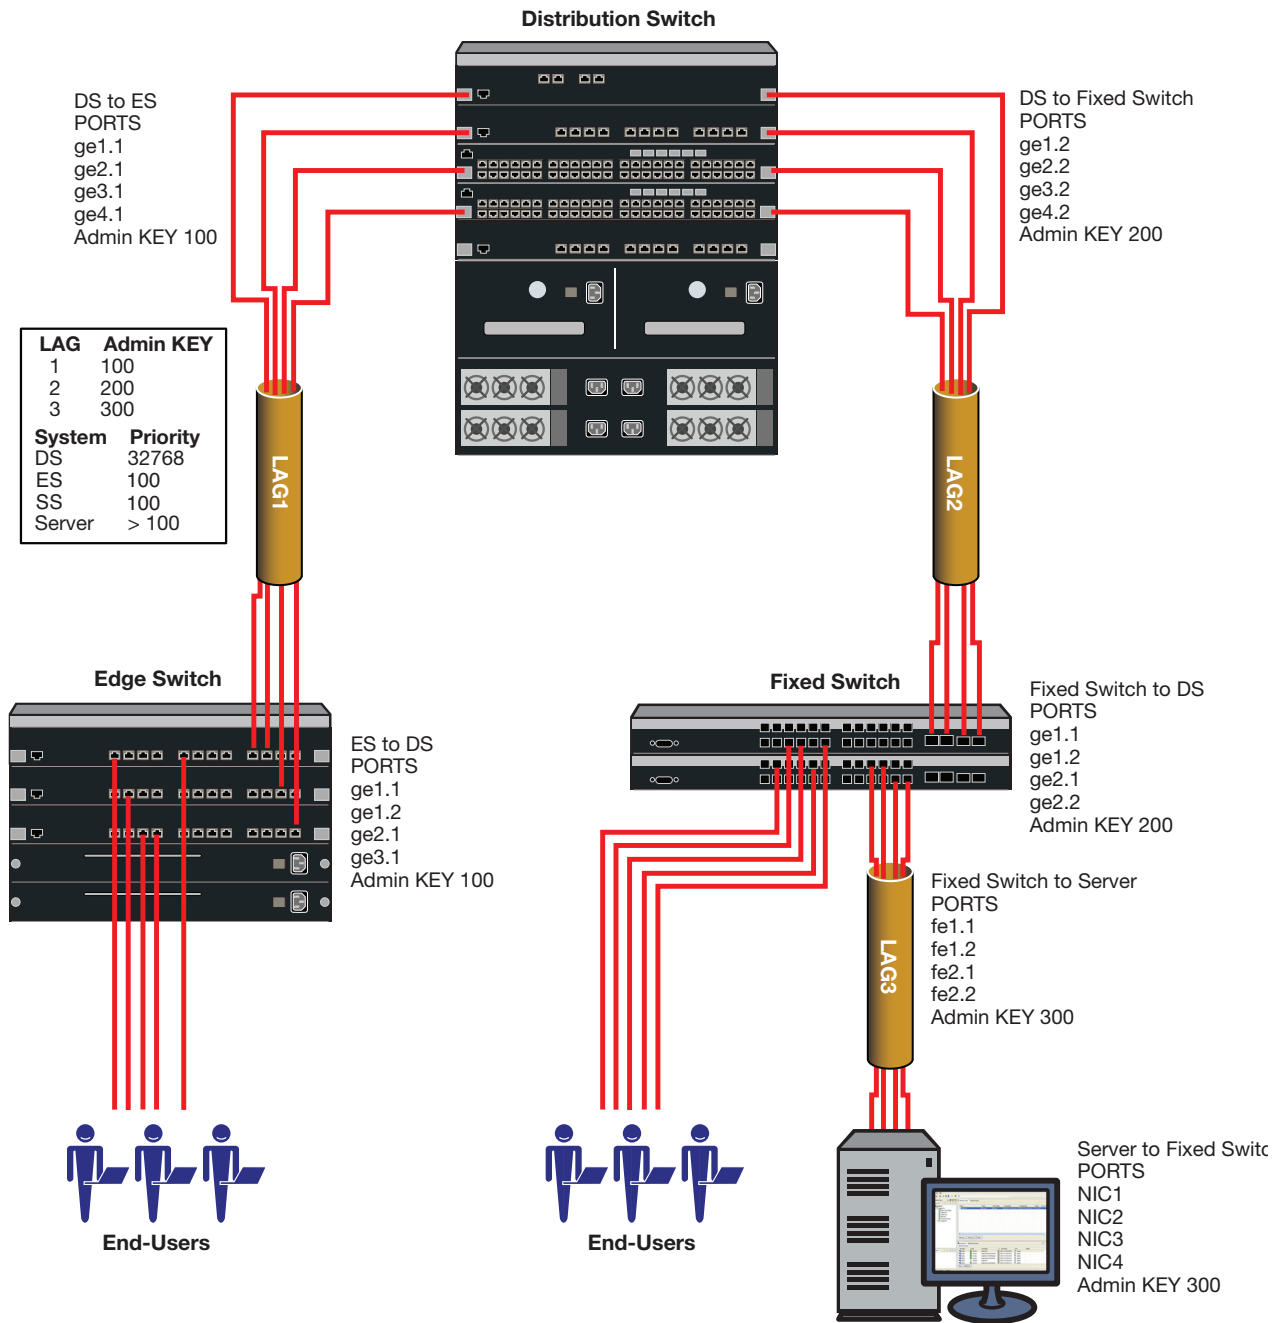
- An example of link aggregations between multiple devices
- An example of link aggregation when a LAG contains physical ports with different speeds

### Link Aggregation Configuration Example 1

This example provides a link aggregation configuration example that includes an edge switch, a distribution switch, and two Fixed Switches that will aggregate both end-users at the edge and the data from a local server.

See [Figure 21-3](#) on page 12 for an illustration of this example, including port, key, and system priority assignments.

**Figure 21-3 Example 1 Multiple Device Configuration**



Three LAGs are created for the example:

- LAG 1 provides an uplink aggregate of four 1Gb ports for the edge switch devices to the distribution switch.
- LAG2 provides an uplink aggregate of four 1Gb ports for the Fixed Switches to the distribution switch for both the end-user and server data flows.
- LAG3 provides an aggregate of four 100Mbps ports between the Fixed Switches and the server.

Each LAG consists of four ports. The primary goal of the aggregates in this example is to provide link and slot redundancy for the affected data streams. With that in mind, LAG members are

spread between available system slots. Four out of the five distribution switch available slots are used providing complete redundancy at the distribution switch. All three slots are used in the edge switch. The four ports from the server to the Fixed Switches and the Fixed Switches to the distribution switch are evenly split between the two Fixed Switches.

For this example we will manually configure the LAGs that will form and prevent any other LAGs from forming. Because we have specific port to LAG goals in mind, the first thing we want to do on each device is to ensure that LAGs form only where we configure them. Since the admin key for the LAG and its associated ports must agree for the LAG to form, an easy way to ensure that LAGs do not automatically form is to set the admin key for all LAGS on all devices to a non-default value. The physical ports will initially retain admin key defaults. In our example, the admin keys for all LAGs are set to the highest configurable value of 65535.

Both physical port and LAG admin keys will be set as shown in [Table 21-7](#) to ensure that the LAGs form only for the desired ports.

**Table 21-7 LAG and Physical Port Admin Key Assignments**

Device	LAG	LAG Admin Key	Physical Port	Physical Port Admin Key
Distribution Switch	1	100	ge.1.1	100
			ge.2.1	100
			ge.3.1	100
			ge.4.1	100
	2	200	ge.1.2	200
			ge.2.2	200
			ge.3.2	200
			ge.4.2	200
Edge Switch	1	100	ge.1.1	100
			ge.1.2	100
			ge.2.1	100
			ge.3.1	100
Fixed Switch	2	200	ge.1.1	200
			ge.1.2	200
			ge.2.1	200
			ge.2.2	200
	3	300	ge.1.1	300
			ge.1.2	300
			ge.2.1	300
			ge.2.2	300
Server	3	300	NIC1 ETH	300
			NIC2 ETH	300
			NIC3 ETH	300
			NIC4 ETH	300

Which device determines port selection for the LAG is an optional consideration. If system priorities remain at the default value, the lowest MAC address device determines port selection for the LAG. For purposes of this example, we will set the system priority of the edge switch to 100 to ensure it will control port selection for LAG1, instead of the distribution switch. The Fixed Switch system priority will be set to 100 to ensure it will control port selection for LAG2, instead of the distribution switch. For the Fixed Switch to control port selection for LAG3 requires that you ensure that the server has a system priority higher than 100.

Each LAG in our example is made up of physical ports of the same speed, so there is no need to set the port priority to a non-default value. The only port value to be changed is the admin key for each physical port and each LAG. These modifications are detailed in [Table 21-7](#) on page 21-13.

Given that the intent of the example is to have three LAGs of 4 ports each, there is no need to enable the single port LAG feature. Once the LAGs initiate, they will persist across resets. Should only a single port be active after a reset, the LAG will form regardless of the single port LAG feature setting.

The output algorithm defaults to selecting the output port based upon the destination and source IP address. This setting will not be changed in our example. In any case, note that the Fixed Switch does not support the output algorithm feature.

## Configuring the Distribution Switch

The first thing we want to do is set the admin key for all LAGs to the non-default value of 65535 so that no LAGs will automatically form:

```
7100 Chassis(rw)->set lacp aadminkey lag.0.* 65535
```

LAGs 1 and 2 will form on the distribution switch so we need to set the admin keys for these LAGs:

```
7100 Chassis(rw)->set lacp aadminkey lag.0.1 100
7100 Chassis(rw)->set lacp aadminkey lag.0.2 200
```

We next want to enable the port active state and set the admin keys for the distribution switch physical ports:

```
7100 Chassis(rw)->set port lacp port ge.1.1 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port ge.2.1 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port ge.3.1 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port ge.4.1 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port ge.1.2 aadminkey 200 enable
7100 Chassis(rw)->set port lacp port ge.2.2 aadminkey 200 enable
7100 Chassis(rw)->set port lacp port ge.3.2 aadminkey 200 enable
7100 Chassis(rw)->set port lacp port ge.4.2 aadminkey 200 enable
```

Because we want the edge switch and the Fixed Switch to be in charge of port selection, the system priority for the distribution switch will be left at the default value of 32768.

## Configuring the Edge Switch

The first thing we want to do is set the admin key for all LAGs to the non-default value of 65535 so that no LAGs will automatically form:

```
7100 Chassis(rw)->set lacp aadminkey lag.0.* 65535
```

LAG 1 will form on the edge switch so we need to set the admin key for this LAG:

```
7100 Chassis(rw)->set lacp aadminkey lag.0.1 100
```

We next want to enable the port active state and set the admin keys for the edge switch physical ports:

```
7100 Chassis(rw)->set port lacp port ge.1.1 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port ge.1.2 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port ge.2.1 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port ge.3.1 aadminkey 100 enable
```

Next we want to change the system priority for the edge switch so that it will be in charge of port selection on LAG1:

```
7100 Chassis(rw)->set lacp asyspri 100
```

## Configuring the Fixed Switch

The first thing we want to do is set the admin key for all LAGs to the non-default value of 65535 so that no LAGs will automatically form:

```
FixedSwitch(rw)->set lacp aadminkey lag.0.* 65535
```

LAGs 2 and 3 will form on the Fixed Switch so we need to set the admin key for this LAG:

```
FixedSwitch(rw)->set lacp aadminkey lag.0.2 200
FixedSwitch(rw)->set lacp aadminkey lag.0.3 300
```

We next want to enable the port active state and set the admin keys for the Fixed Switch physical ports:

```
FixedSwitch(rw)->set port lacp port ge.1.1 aadminkey 200 enable
FixedSwitch(rw)->set port lacp port ge.1.2 aadminkey 200 enable
FixedSwitch(rw)->set port lacp port ge.2.1 aadminkey 200 enable
FixedSwitch(rw)->set port lacp port ge.2.2 aadminkey 200 enable
FixedSwitch(rw)->set port lacp port ge.1.1 aadminkey 300 enable
FixedSwitch(rw)->set port lacp port ge.1.2 aadminkey 300 enable
FixedSwitch(rw)->set port lacp port ge.2.1 aadminkey 300 enable
FixedSwitch(rw)->set port lacp port ge.2.2 aadminkey 300 enable
```

Next we want to change the system priority for the Fixed Switch so that it will be in charge of port selection on LAGs 2 and 3:

```
FixedSwitch(rw)->set lacp asyspri 100
```

## Configuring the Server

Configuring link aggregation on the server is dependent upon the installed LACP application. There are three aspects to link aggregation on the server you must ensure for this example:

- The admin key for LAG3 must be set to 300
- The admin keys for each NIC port must be set to 300
- The system priority for the server must be set greater than 100 to ensure that the Fixed Switch will control port selection

This completes the example 1 configuration.

## Link Aggregation Configuration Example 2

It is unlikely that you will run out of LAG resources for most link aggregation configurations, but it is possible. See [Table 21-3](#) on page 21-8 for a listing of LAG support for your system. Should you run out of LAG resources, excess aggregatable ports are placed in standby mode.

Making use of the port priority parameter, this example shows how you can ensure the order in which aggregatable ports form a LAG and are moved to the attached state. In configuration example 2, two uplink LAGs will be manually configured between two edge switch chassis. The first LAG consists of two 1 Gb ports. The second LAG consists of eight 100 Mbps ports. In this example we will ensure that the two 1Gb port LAG forms before the eight 100 Mbps port LAG.



See [Figure 21-4](#) on page 21-17 for an illustration of this example, including port, key and port priority assignments.

The LAG configuration will ensure that the two 1Gb ports attach to the first available LAG (LAG1). The eight 100Mbps ports will then attach to the second available LAG (LAG2)

Which device determines port selection for the LAG is an optional consideration. For this example, system priorities are not modified, the lowest MAC address device will determine port selection for the LAG.


There are two physical port speeds in our example, 100Mbps and 1Gb. A LAG only moves ports of the same speed to the attached state. Selecting the ports to move to attached state is based upon the lowest port priority. If port priorities are the same, the lowest port number breaks the tie. For our example, we want to ensure that the 1Gb ports are moved to the attached state for LAG1. Port priority for 1Gb ports is set to 100. Port priority for 100Mbps ports is left at the default value of 32768.

The admin key for each 100 Mbps to 1Gb physical port link and LAG in the example is set to 100, and for each 1Gb to 1Gb physical port link and LAG is set to 200. This ensures that LAGs will form for each set of ports.

For this example we will allow single port LAGs to form. The single port LAG feature will be set to enabled for both devices.

The output algorithm defaults to selecting the output port based upon the destination and source IP address. This setting will not be changed in our example.

Figure 21-4 Example 2 Configuration



## Configuring the Edge Switch

For this example, we want LAGs to form wherever they can so we will not change the default admin key setting for all LAGs as we did in the multiple device example. Because we want LAG1 and LAG2, as described for this example, to form for specific ports, we set separate admin keys for these LAGs (LAG1 to 200 and LAG2 to 100):

```
7100 Chassis(rw)->set lacp aadminkey lag.0.1 200
7100 Chassis(rw)->set lacp aadminkey lag.0.2 100
```

We next want to enable the port active state and set the admin keys for the edge switch physical ports associated with LAG1 and LAG2:

```
7100 Chassis(rw)->set port lacp port ge.2.1 aadminkey 200 enable
7100 Chassis(rw)->set port lacp port ge.3.1 aadminkey 200 enable
7100 Chassis(rw)->set port lacp port fe.1.1 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port fe.1.2 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port fe.1.3 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port fe.1.4 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port fe.1.5 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port fe.1.6 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port fe.1.7 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port fe.1.8 aadminkey 100 enable
```

System priority determines which device will be in charge of port selection. This is an optional consideration. For this example we will leave system priority at the default value and allow the device with the lowest MAC address to determine port selection.

Port priority determines which aggregatable ports available for a LAG are moved to the attached state when different speed physical ports form a LAG. For this example we want to ensure that the 1Gb ports move to the attached state for LAG1. We will set the port priority to 100 for the 1Gb actor ports should this device be in charge of selecting ports to move to the attached state:

```
7100 Chassis(rw)->set port lacp port ge.2.1 aporti 100
7100 Chassis(rw)->set port lacp port ge.3.1 aporti 100
```

We next enable single port LAGs on this device:

```
7100 Chassis(rw)->set lacp singleportlag enable
```

## Configuring the Upstream Switch

For this example, we want LAGs to form wherever they can so we will not change the default admin key setting for all LAGs as we did in the multiple device example. Because we want LAG1 and LAG2, as described for this example, to form for specific ports, we set the admin key of LAG 1 to 200 and the admin key of LAG 2 to 100:

```
7100 Chassis(rw)->set lacp aadminkey lag.0.1 200
7100 Chassis(rw)->set lacp aadminkey lag.0.2 100
```

We next want to enable the port active state and set the admin keys for the upstream switch physical ports associated with LAG 1 and LAG 2:

```
7100 Chassis(rw)->set port lacp port ge.2.1 aadminkey 200 enable
7100 Chassis(rw)->set port lacp port ge.3.1 aadminkey 200 enable
7100 Chassis(rw)->set port lacp port ge.1.1 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port ge.1.2 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port ge.1.3 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port ge.1.4 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port ge.2.1 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port ge.2.2 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port ge.2.3 aadminkey 100 enable
7100 Chassis(rw)->set port lacp port ge.2.4 aadminkey 100 enable
```

System priority determines which device will be in charge of port selection. This is an optional consideration. For this example we will leave system priority at the default value and allow the device with the lowest MAC address to determine port selection.

Port priority determines which aggregatable ports available for a LAG are moved to the attached state when different speed physical ports form a LAG. For this example we want to ensure that the 1Gb ports move to the attached state for LAG1. We will set the port priority to 100 for the 1Gb actor ports should this device be in charge of selecting ports to move to the attached state:

```
7100 Chassis(rw)->set port lacp port ge.2.1 aporti 100
7100 Chassis(rw)->set port lacp port ge.3.1 aporti 100
```

We next enable single port LAGs on this device:

```
7100 Chassis(rw) ->set lacp singleportlag enable
```

This completes the example 2 configuration.

## Terms and Definitions

[Table 21-8](#) lists terms and definitions used in this link aggregation configuration discussion.

**Table 21-8 Link Aggregation Configuration Terms and Definitions**

Term	Definition
Aggregator	Virtual port that controls link aggregation for underlying physical ports. Each device provides aggregator ports, which are designated in the CLI as <b>lag.0.1</b> through <b>lag.0.x</b> (depending upon the device, see <a href="#">Table 21-3</a> on page 21-8 for LAG resources available on your device).
LAG	Link Aggregation Group. Once underlying physical ports (i.e.; <b>ge.x.x</b> ) are associated with an aggregator port, the resulting aggregation will be represented as one LAG with a <b>lag.x.x</b> port designation.
LACPDU	Link Aggregation Control Protocol Data Unit. The protocol exchanges aggregation state/mode information by way of a port's actor and partner operational states. LACPDUs sent by the first party (the actor) convey to the second party (the actor's protocol partner) what the actor knows, both about its own state and that of its partner.
Actor and Partner	An actor is the local device sending LACPDUs. Its protocol partner is the device on the other end of the link aggregation. Each maintains current status of the other via LACPDUs containing information about their ports' LACP status and operational state.
Admin Key	Value assigned to aggregator ports and physical ports that are candidates for joining a LAG. The LACP implementation uses this value to determine which underlying physical ports are capable of aggregating by comparing keys. Aggregator ports allow only underlying ports with admin keys that match the aggregator to join their LAG.
Port Priority	Port priority determines which physical ports are moved to the attached state when physical ports of differing speeds form a LAG. Port priority also determines which ports will join a LAG when the number of supported ports for a LAG is exceeded.
System Priority	Value used to build a LAG ID, which determines aggregation precedence. If there are two partner devices competing for the same aggregator, LACP compares the LAG IDs for each grouping of ports. The LAG with the lower LAG ID is given precedence and will be allowed to use the aggregator.



## Policy Configuration

This document describes the Extreme Networks policy feature and its configuration on Extreme Networks 7100-Series devices.

For information about...	Refer to page...
<a href="#">Using Policy in Your Network</a>	<a href="#">22-1</a>
<a href="#">Implementing Policy</a>	<a href="#">22-2</a>
<a href="#">Policy Overview</a>	<a href="#">22-2</a>
<a href="#">Configuring Policy</a>	<a href="#">22-10</a>
<a href="#">Policy Configuration Example</a>	<a href="#">22-15</a>
<a href="#">Terms and Definitions</a>	<a href="#">22-25</a>

### Using Policy in Your Network

Policy is a component of Secure Networks that provides for the configuration of role-based profiles for securing and provisioning network resources based upon the role the user or device plays within the enterprise. By first defining the user or device role, network resources can be granularly tailored to a specific user, system, service, or port-based context by configuring and assigning rules to the policy role. A policy role can be configured for any combination of Class of Service, VLAN assignment, classification rule precedence, logging, accounting, or default behavior based upon L2, L3, and L4 packet fields. Hybrid authentication allows either policy or dynamic VLAN assignment, or both, to be applied through RADIUS authorization.

The three primary benefits of using Extreme Networks Secure Networks policy in your network are provisioning and control of network resources, security, and centralized operational efficiency using the Extreme Networks NetSight Policy Manager.

Policy provides for the provisioning and control of network resources by creating policy roles that allow you to determine network provisioning and control at the appropriate network layer, for a given user or device. With a role defined, rules can be created based upon up to 15 traffic classification types for traffic drop or forwarding, with classification support determined by the system resource allocation profile configured. A Class of Service (CoS) can be associated with each role for purposes of setting priority, forwarding queue, rate limiting, and rate shaping.

Security can be enhanced by allowing only intended users and devices access to network protocols and capabilities. Some examples are:

- Ensuring that only approved stations can use SNMP, preventing unauthorized stations from viewing, reading, and writing network management information
- Preventing edge clients from attaching network services that are appropriately restricted to data centers and managed by the enterprise IT organization such as DHCP and DNS services

- Identifying and restricting routing to legitimate routing IP addresses to prevent DoS, spoofing, data integrity and other routing related security issues
- Ensuring that FTP/TFTP file transfers and firmware upgrades only originate from authorized file and configuration management servers
- Preventing clients from using legacy protocols such as IPX, AppleTalk, and DECnet that should no longer be running on your network

Extreme Networks NetSight Policy Manager provides a centralized point and click configuration, and one click pushing of defined policy out to all network elements. Use the Extreme Networks NetSight Policy Manager for ease of initial configuration and response to security and provisioning issues that may come up during real-time network operation.

## Implementing Policy

To implement policy:

- Identify the roles of users and devices in your organization that access the network
- Create a policy role for each identified user role
- Associate classification rules and administrative profiles with each policy role
- Optionally, configure a class of service and associate it directly with the policy role or through a classification rule
- Optionally, change the system resource allocation profile
- Optionally, set device response to invalid policy

## Policy Overview

### Introduction

This section provides an overview of policy configuration. Policy is implemented on an Extreme Networks platform by associating users and devices in the network with defined enterprise roles (such as sales, engineering, or administration) that are configured in a policy role. The policy role is associated with rules that define how network resources will be provisioned and controlled for role members, as well as how security will be applied to the role member. An administrative profile associates a specific role member traffic classification with a policy role.



**Note:** In a CLI configuration context, the policy role is configured within a policy profile using the **set policy profile** command. throughout this discussion, policy role and policy profile mean the same thing.

### The Extreme Networks NetSight Policy Manager

Extreme Networks NetSight Policy Manager is a management GUI that automates the definition and enforcement of network-wide policy rules. It eliminates the need to configure policies on a device-by-device basis using complex CLI commands. The Policy Manager's GUI provides ease of classification rule and policy role creation, because you only define policies once using an easy to understand point and click GUI— and regardless of the number of moves, adds or changes to the policy role, Policy Manager automatically enforces roles on Extreme Networks security-enabled infrastructure devices.

This document presents policy configuration from the perspective of the CLI. Though it is possible to configure policy from the CLI, CLI policy configuration in even a small network can be

prohibitively complex from an operational point of view. It is highly recommended that policy configuration be performed using the NetSight Policy Manager. The NetSight Policy Manager provides:

- Ease of rule and policy role creation
- The ability to store and retrieve roles and policies
- The ability, with a single click, to enforce policy across multiple devices

The official Policy Manager documentation is accessed using online help from within the application. This online documentation completely covers the configuration of policy in a Policy Manager context. For access to the Policy Manager data sheet or to setup a demo of the product, see <http://www.extremenetworks.com/product/netsight/>.

## Understanding Roles in a Secure Network

The capacity to define roles is directly derived from the ability of the Extreme Networks 7100-Series, stackable, and standalone devices to isolate packet flows by inspecting Layer 2, Layer 3, and Layer 4 packet fields while maintaining line rate. This capability allows for the granular application of a policy to a:

- Specific user (MAC, IP address or interface)
- Group of users (masked MAC or IP address)
- System (IP address)
- Service (such as TCP or UDP)
- Port (physical or application)

Because users, devices, and applications are all identifiable within a flow, a network administrator has the capacity to define and control network access and usage by the actual role the user or device plays in the network. The nature of the security challenge, application access, or amount of network resource required by a given attached user or device, is very much dependent upon the “role” that user or device plays in the enterprise. Defining and applying each role assures that network access and resource usage align with the security requirements, network capabilities, and legitimate user needs as defined by the network administrator.

### The Policy Role

A role, such as sales, admin, or engineering, is first identified and defined in the abstract as the basis for configuring a policy role. Once a role is defined, a policy role is configured and applied to the appropriate context using a set of rules that can control and prioritize various types of network traffic. The rules that make up a policy role contain both classification definitions and actions to be enforced when a classification is matched. Classifications include Layer 2, Layer 3, and Layer 4 packet fields. Policy actions that can be enforced include VLAN assignment, filtering, inbound rate limiting, outbound rate shaping, and priority class mapping.

## Policy Roles

### Defining a Policy Role

The policy role is a container that holds all aspects of policy configuration for a specific role. Policy roles are identified by a numeric profile-index value between **1** and the maximum number of roles supported on the platform. Please see your device’s firmware release notes for the maximum number of roles supported. Policy roles are configured using the **set policy profile** command. Policy configuration is either directly specified with the **set policy profile** command or is



associated with the role by specifying the profile-index value within the command syntax where the given policy option is configured. For example, when configuring a policy mappable entry using the **set policy mappable** command (see “[VLAN-to-Policy Mapping](#)” on page 22-5), the command syntax requires that you identify the policy role the mappable entry will be associated with, by specifying the profile-index value.

When modifying an existing policy role the default behavior is to replace the existing role with the new policy role configuration. Use the **append** option to limit the change to the existing policy role to the options specified in the entered command.

A policy role can also be identified by a text name of between 1 and 64 characters. This name value is used by the RADIUS filter-ID attribute to identify the policy role to be applied by the switch with a successful authentication.

The following example creates a policy profile with a profile-index value of **1** and a profile name, **student**, to be used by the RADIUS filter-ID functionality (tci-overwrite is always enabled):

```
7100 Chassis(rw)->set policy profile 1 student
```

## Setting a Default VLAN for this Role

A default VLAN can be configured for a policy role. A default VLAN will only be used when either a VLAN is not specifically assigned by a classification rule or all policy role classification rules are missed. To configure a default VLAN, enable **pvid-status** and specify the port VLAN to be used. **pvid-status** is disabled by default.



**Note:** Extreme Networks supports the assignment of port VLAN-IDs 1 - 4094. VLAN-IDs 0 and 4095 can not be assigned as port VLAN-IDs, but do have special meanings within a policy context and can be assigned to the **pvid** parameter (See “[VLAN Support on Extreme Networks S-Series Switches](#)” on page 24-6 for further information on these two VLAN-IDs). Within a policy context:

- **0** - Specifies an explicit deny all
- **4095** - Specifies an explicit permit all

The following example creates a policy profile with a profile-index value of **1** and associates with it a default VLAN with an ID of **2**.

```
7100 Chassis(rw)->set policy profile 1 pvid-status enable pvid 2
```

## Assigning a Class of Service to this Role

How a packet is treated as it transits the link can be configured in the Class of Service (CoS). It is through a CoS that Quality of Service (QoS) is implemented. A CoS can be configured for the following values:

- 802.1p priority
- IP Type of Service (ToS) rewrite value
- Priority Transmit Queue (TxQ) along with a forwarding behavior
- Inbound rate limiter per transmit queue
- Outbound rate shaper per transmit queue

CoS configurations are identified by a numeric value between 0 - 255. 0 - 7 are fixed 802.1p CoS configurations. CoS configurations 8 - 255 are user configurable. Policy uses the **cos** option followed by the CoS configuration ID value to associate a CoS with a policy role.

See [Chapter 40, Quality of Service \(QoS\) Configuration](#) for a complete discussion of QoS configuration.

The following example creates a policy profile with a profile-index value of **1** and associates with the profile a user configured CoS **8**:

```
7100 Chassis(rw)->set policy profile 1 cos-status enable cos 8
```

## Adding Tagged, Untagged, and Forbidden Ports to the VLAN Egress Lists

The VLAN Egress list contains a list of ports that a frame for this VLAN can exit. Specified ports are automatically assigned to the VLAN egress list for this policy role as tagged, untagged, or forbidden. Ports are added to the VLAN egress list using the **egress-vlans**, **forbidden-vlans**, and **untagged-vlans** options of the **set policy profile** command.

## VLAN-to-Policy Mapping

VLAN-to-Policy mapping provides for the manual configuration of a VLAN-to-Policy association that creates a policy mappable entry between the specified VLAN and the specified policy role. A policy mappable holds the VLAN-to-Policy mappings. When an incoming tagged VLAN packet is seen by the switch, a lookup of the policy mappable determines whether a VLAN-to-policy mapping exists. If the mapping exists, the associated policy is applied to this packet.

This feature can be used at the distribution layer in environments where non-policy capable edge switches are deployed and there is no possibility of applying Extreme Networks policy at the edge. Tagged frames received at the distribution layer interface for a VLAN with an entry in the policy mappable will have the associated policy applied to the frame.

Use the **set policy mappable** command specifying a single VLAN ID or range of IDs and the policy profile-index to create a policy mappable entry.

The following example creates a policy mappable entry associating VLAN **100** and policy profile **10**:

```
7100 Chassis(rw)->set policy mappable 100 10
```

## Applying Policy Using the RADIUS Response Attributes

If an authentication method that requires communication with an authentication server is configured for a user, the RADIUS filter-ID attribute can be used to dynamically assign a policy role to the authenticating user. Supported RADIUS attributes are sent to the switch in the RADIUS access-accept message. The RADIUS filter-ID can also be applied in hybrid authentication mode. Hybrid authentication mode determines how the RADIUS filter-ID and the three RFC 3580 VLAN tunnel attributes (VLAN Authorization), when either or all are included in the RADIUS access-accept message, will be handled by the switch. The three VLAN tunnel attributes define the base VLAN-ID to be applied to the user. In either case, conflict resolution between RADIUS attributes is provided by the mappable response feature.



**Note:** VLAN-to-policy mapping to mappable response configuration behavior is as follows:

- If the RADIUS response is set to **policy**, any VLAN-to-policy mappable configuration is ignored for all platforms.
- If the RADIUS response is set to **tunnel**, VLAN-to-policy mapping can occur on an 7100-Series platform.
- If the RADIUS response is set to **both** and both the filter-ID and tunnel attributes are present, VLAN-to-policy mapping configuration is ignored.

See the “[Policy Mappable Response](#)” on page 43-13 for a detailed RADIUS response discussion.

Please see “[Configuring RADIUS](#)” on page 43-28 for a discussion of RADIUS configuration, the RADIUS filter-ID, and VLAN authorization.

Use the **policy** option of the **set policy mactable response** command to configure the switch to dynamically assign a policy using the RADIUS filter-ID in the RADIUS response message.

The following example specifies that the RADIUS filter-ID, if it is present in the RADIUS response message when a user authenticates, should be used to apply the specified policy to the user:

```
7100 Chassis(rw)->set policy mactable response policy
```

## Applying Policy Using Hybrid Authentication Mode

Hybrid authentication is an authentication capability that allows the switch to use both the filter-ID and tunnel attributes in the RADIUS response message to determine how to treat the authenticating user.

Hybrid authentication is configured by specifying the **both** option in the **set policy mactable response** command. The **both** option:

- Applies the VLAN tunnel attributes if they exist and the filter-ID attribute does not
- Applies the filter-ID attribute if it exists and the VLAN tunnel attributes do not
- Applies both the filter-ID and the VLAN tunnel attributes if all attributes exist

If all attributes exist, the following rules apply:

- The policy role will be enforced, with the exception that any port PVID specified in the role will be replaced with the VLAN tunnel attributes
- The policy map is ignored because the policy role is explicitly assigned
- VLAN classification rules are assigned as defined by the policy role

**vlanauthorization** must be enabled or the VLAN tunnel attributes are ignored and the default VLAN is used. Please see [“Configuring VLAN Authorization”](#) on page 43-27 for a complete VLAN Authorization discussion.

Hybrid Mode support eliminates the dependency of VLAN assignment based on roles. As a result, VLANs can be assigned via the tunnel-private-group-ID, as defined per RFC3580, while assigning roles via the filter-ID. This separation gives administrators more flexibility to segment their networks for efficiency beyond the role limits associated with the B3, C3, and G3 platforms.

The following example specifies that either or both the vlan-tunnel and filter-ID attributes can be included in the RADIUS response message:

```
7100 Chassis(rw)->set policy mactable response both
```

## Device Response to Invalid Policy

The action that the device should take when asked to apply an invalid or unknown policy can be specified. The available actions are:

- Ignore the result and search for the next policy assignment rule. If all rules are missed, the default policy is applied.
- Block traffic
- Forward traffic as if no policy has been assigned using 802.1D/Q rules

Use the **set policy invalid action** command to specify a default action to take when asked to apply an invalid or unknown policy.

The following example specifies that an attempt to apply an invalid or unknown policy should be ignored:

```
7100 Chassis(rw)->set policy invalid action default-policy
```

## Classification Rules

Classification rules associate specific traffic classifications or policy behaviors with the policy role. There are two aspects of classification rule configuration:

- The association of a traffic classification with a policy role by assigning the traffic classification to an administrative profile.
- The assignment of policy rules that define desired policy behaviors for the specified traffic classification type.

Both the administrative profile and policy rules are associated with the policy role by specifying the **admin-pid** option, in the case of an administrative profile, or a **profile-index** value, in the case of the policy rule. Administrative profiles and policy rules are configured using the **set policy rule** command.

The administrative profile assigns a traffic classification to a policy role by using the **admin-profile** option of the **set policy rule** command.

Policy rules are based on traffic classifications. [Table 22-1](#) on page 22-7 provides the supported policy rule traffic classification command options and definitions.

A detailed discussion of supported traffic classifications is available in the “Traffic Classification Rules” section of the NetSight Policy Manager online help.

**Table 22-1 Administrative Policy and Policy Rule Traffic Classifications**

Traffic Classification	Description	Attribute ID
<b>macsource</b>	Classifies based on MAC source address.	<b>1</b>
<b>macdest</b>	Classifies based on MAC destination address.	<b>2</b>
<b>ip6dest</b>	Classifies based on IPv6 destination address.	<b>10</b>
<b>ipsourcesocket</b>	Classifies based on source IP address.	<b>12</b>
<b>ipdestsocket</b>	Classifies based on destination IP address.	<b>13</b>
<b>ip frag</b>	Classifies based on IP fragmentation value.	<b>14</b>
<b>udpsourceportip</b>	Classifies based on UDP source port.	<b>15</b>
<b>udpdestportip</b>	Classifies based on UDP destination port.	<b>16</b>
<b>tcpsourceportip</b>	Classifies based on TCP source port.	<b>17</b>
<b>tcpdestportip</b>	Classifies based on TCP destination port.	<b>18</b>
<b>ipttl</b>	Classifies based on Time-To-Live (TTL).	<b>20</b>
<b>iptos</b>	Classifies based on Type of Service field in IP packet.	<b>21</b>
<b>ipproto</b>	Classifies based on protocol field in IP packet.	<b>22</b>
<b>ether</b>	Classifies based on type field in Ethernet II packet.	<b>25</b>
<b>port</b>	Classifies based on port-string.	<b>31</b>

A data value is associated with most traffic classifications to identify the specific network element for that classification. For data value and associated mask details, see the “Valid Values for Policy Classification Rules” table in the **set policy rule** command discussion of the command reference guide for your platform.

## Specifying Storage Type

Specifying the storage type for a rule entry is supported. Storage types are **volatile** and **non-volatile**. Volatile storage does not persist after a reset of the device. Non-volatile storage does persist after a reset of the device. Use the **storage-type** option to specify the desired storage type for this policy rule entry.

## Forward and Drop

Packets for this entry can be either forwarded or dropped for this traffic classification using the **forward** and **drop** policy rule options.

## Allowed Traffic Rule-Type on a Port

Use the **show policy allowed-type** command to display a table of the current allowed and disallowed traffic rule-types for the specified port(s).

See [Table 22-1](#) on page 22-7 for a listing of supported traffic classification rule-types.

## Quality of Service in a Policy Rules Context

Quality of Service (QoS) can be specified directly in a policy role as stated in “[Assigning a Class of Service to this Role](#)” on page 22-4. A CoS can also be applied to a policy rule. The CoS specified at the policy role level is the default and is only used if no rule is triggered. Therefore, if a CoS is applied to both the policy role and a policy rule, the CoS specified in the policy rule takes precedence over the CoS in the policy role for the traffic classification context specified in the policy rule. As stated in the policy role discussion, CoS configuration details are beyond the scope of this document. See [Chapter 40, Quality of Service \(QoS\) Configuration](#) for a complete discussion of QoS configuration.

The following example applies CoS 8 to profile-index 1 for port ge.1.1:

```
7100 Chassis(rw)->set policy rule 1 port ge.1.1 port-string ge.1.1 cos 8
```

## Blocking Non-Edge Protocols at the Edge Network Layer

Edge clients should be prevented from acting as servers for a number of IP services. If non-edge IP services accidentally or maliciously attach to the edge of the network, they are capable of disrupting network operation. IP services should only be allowed where and when your network design requires. This section identifies ten IP Services you should consider blocking at the edge unless allowing them is part of your network architecture.

**Table 22-2 Non-Edge Protocols**

Protocol	Policy Effect
<b>DHCP Server Protocol</b>	Every network needs DHCP. Automatically mitigate the accidental or malicious connection of a DHCP server to the edge of your network to prevent DoS or data integrity issues, by blocking DHCP on the source port for this device.
<b>DNS Server Protocol</b>	DNS is critical to network operations. Automatically protect your name servers from malicious attack or unauthorized spoofing and redirection, by blocking DNS on the source port for this device.
<b>Routing Topology Protocols</b>	RIP, OSPF, and BGP topology protocols should only originate from authorized router connection points to ensure reliable network operations.

**Table 22-2 Non-Edge Protocols (continued)**

Protocol	Policy Effect
<b>Router Source MAC and Router Source IP Address</b>	Routers and default gateways should not be moving around your network without approved change processes being authorized. Prevent DoS, spoofing, data integrity and other router security issues by blocking router source MAC and router source IP addresses at the edge.
<b>SMTP/POP Server Protocols</b>	Prevent data theft and worm propagation by blocking SMTP at the edge.
<b>SNMP Protocol</b>	Only approved management stations or management data collection points need to be speaking SNMP. Prevent unauthorized users from using SNMP to view, read, or write management information.
<b>FTP and TFTP Server Protocols</b>	Ensure file transfers and firmware upgrades are only originating from authorized file and configuration management servers.
<b>Web Server Protocol</b>	Stop malicious proxies and application-layer attacks by ensuring only the right Web servers can connect from the right location at the right time, by blocking HTTP on the source port for this device.
<b>Legacy Protocols</b>	If IPX, AppleTalk, DECnet or other protocols should no longer be running on your network, prevent clients from using them. Some organizations even take the approach that unless a protocol is specifically allowed, all others are denied.

## System Resource Allocation Profile

The system resource allocation profile determines the traffic classifications and ACL types supported for the system and the number of rules supported for each traffic classification and ACL type. There are currently two supported profiles: **default** and **router1**.

Policy traffic classifications are broken into four sets including MAC rules, IPv6 rules, IPv4 rules, and L2 rules configured using the **set policy rule** command. Traffic classifications are applied to ACLs as described in [Chapter 52, Access Control List Configuration](#).

The 7100-Series supports up to 512 admin policy rules based upon the macsource and port rule classifications and 768 non-admin policy rules. Within the non-admin policy rule support, limits are placed on the number of rules within a policy traffic classification set. Non-admin policy rules belonging to each set are defined in [Table 22-3](#) on page 22-10.

See the release notes that come with your firmware for the supported number of traffic classification and ACL rules supported for each system resource allocation profile.

**Table 22-3 Policy Traffic Classification Sets**

MAC Rules	IPv6 Rules	IPv4 Rules	L2 Rules
<ul style="list-style-type: none"> <li>• macsource</li> <li>• macdest</li> </ul>	<ul style="list-style-type: none"> <li>• ipv6dest</li> </ul>	<ul style="list-style-type: none"> <li>• ipsourcesocket</li> <li>• ipdestsocket</li> <li>• ipfrag</li> <li>• udpsourceportIP</li> <li>• udpdestportIP</li> <li>• tcpsourceportIP</li> <li>• tcpdestportIP</li> <li>• ipttl</li> <li>• iptos</li> <li>• iptype</li> </ul>	<ul style="list-style-type: none"> <li>• ethertype</li> <li>• port</li> </ul>

Use the **set limits resource-profile** command to set the system resource allocation policy profile for this system.

Use the **clear limits resource-profile** command to reset the system resource allocation policy profile to the **default** profile.

Use the **show limits resource-profile** command to display the current profile setting.

## Policy Capabilities

[Table 22-4](#) provides a listing of policy capabilities.

**Table 22-4 Traffic Classification Based Policy Capabilities**

Traffic Classification	Description
Dynamic PID Assign Rule	The ability to dynamically assign a policy based upon a traffic classification.
Admin PID Assign Rule	The ability to administratively assign a policy based upon a traffic classification.
Deny	The ability to assign a drop traffic rule.
Permit	The ability to assign a forward traffic rule.
CoS Assign Rule	The ability to assign a CoS rule.
Priority	The ability to assign traffic priority using a CoS assignment.
Longest Prefix Rules	The ability to always look at the highest bit mask for an exact traffic classification match.
Invalid Policy Action	The ability to set a drop, forward, or default-policy behavior based upon an invalid action.

## Configuring Policy

This section presents configuration procedures and tables including command description and syntax in the following policy areas: profile, classification, and display.

[Procedure 22-1](#) describes how to configure policy roles and related functionality.

## Procedure 22-1 Configuring Policy Roles

Step	Task	Command(s)
1.	<p>In any command mode, create a policy role.</p> <ul style="list-style-type: none"> <li>• <b>name</b> – (Optional) Specifies a name for this policy profile; used by the filter-ID attribute. This is a string from 1 to 64 characters.</li> <li>• <b>pvid-status</b> – (Optional) Enables or disables PVID override for this policy profile. If all the classification rules associated with this profile are missed, then this parameter, if specified, determines the default VLAN for this profile.</li> <li>• <b>pvid</b> – (Optional) Specifies the PVID to assign to packets, if PVID override is enabled and invoked as the default behavior.</li> <li>• <b>cos-status</b> – (Optional) Enables or disables Class of Service override for this policy profile. If all the classification rules associated with this profile are missed, then this parameter, if specified, determines the default CoS assignment.</li> <li>• <b>cos</b> – (Optional) Specifies a CoS value to assign to packets, if CoS override is enabled and invoked as the default behavior. Valid values are 0 to 255.</li> <li>• <b>egress-vlans</b> – (Optional) Specifies the port to which this policy profile is applied should be added to the egress list of the VLANs defined by egress-vlans. Packets will be formatted as tagged.</li> <li>• <b>forbidden-vlans</b> – (Optional) Specifies the port to which this policy profile is applied should be added as forbidden to the egress list of the VLANs defined by forbidden-vlans. Packets from this port will not be allowed to participate in the listed VLANs.</li> <li>• <b>untagged-vlans</b> – (Optional) Specifies the port to which this policy profile is applied should be added to the egress list of the VLANs defined by untagged-vlans. Packets will be formatted as untagged.</li> <li>• <b>append</b> – (Optional) Appends any egress, forbidden, or untagged specified VLANs to the existing list. If append is not specified, all previous settings for this VLAN list are replaced</li> <li>• <b>clear</b> – (Optional) Clears any egress, forbidden or untagged VLANs specified from the existing list.</li> <li>• <b>tci-overwrite</b> – (Optional) Enables or disables TCI (Tag Control Information) overwrite for this profile. When enabled, rules configured for this profile are allowed to overwrite user priority and other classification information in the VLAN tag's TCI field.</li> </ul>	<pre>set policy profile <i>profile-index</i> [<b>name</b> <i>name</i>] [<b>pvid-status</b> {<b>enable</b>   <b>disable</b>}] [<b>pvid</b> <i>pvid</i>] [<b>cos-status</b> {<b>enable</b>   <b>disable</b>}] [<b>cos</b> <i>cos</i>] [<b>egress-vlans</b> <i>egress-vlans</i>] [<b>forbidden-vlans</b> <i>forbidden-vlans</i>] [<b>untagged-vlans</b> <i>untagged-vlans</i>] [<b>append</b>] [<b>clear</b>] [<b>tci-overwrite</b> {<b>enable</b>   <b>disable</b>}]</pre>



**Procedure 22-1 Configuring Policy Roles (continued)**

Step	Task	Command(s)
2.	<p>(Optional) Assign the action the device will apply to an invalid or unknown policy.</p> <ul style="list-style-type: none"> <li>• <b>default-policy</b> – Instructs the device to ignore this result and search for the next policy assignment rule.</li> <li>• <b>drop</b> – Instructs the device to block traffic.</li> <li>• <b>forward</b> – Instructs the device to forward traffic.</li> </ul>	<b>set policy invalid action</b> { <b>default-policy</b>   <b>drop</b>   <b>forward</b> }
3.	<p>(Optional) Set a policy mappable entry that associates a VLAN with a policy profile. This option is also supported by the B3, C3, and G3 for releases 6.3 and greater.</p>	<b>set policy mappable</b> { <i>vlan-list profile-index</i> }
4.	<p>Optionally, set a policy mappable response.</p> <ul style="list-style-type: none"> <li>• <b>tunnel</b> - Applies the VLAN tunnel attribute.</li> <li>• <b>policy</b> - Applies the policy specified in the filter-ID.</li> <li>• <b>both</b> - Applies either or all the filter-ID and VLAN tunnel attributes or the policy depending upon whether one or both are present.</li> </ul>	<b>set policy mappable response</b> { <b>tunnel</b>   <b>policy</b>   <b>both</b> }

[Procedure 22-2](#) describes how to configure classification rules as an administrative profile or to assign policy rules to a policy role.

**Procedure 22-2 Configuring Classification Rules**

Step	Task	Command(s)
1.	<p>In any command mode, optionally set an administrative profile to assign traffic classifications to a policy role.</p> <p>See <a href="#">Table 22-1</a> on page 22-7 for traffic classification-type descriptions.</p> <p>See the <b>set policy rule</b> command discussion in the command reference guide that comes with your device for traffic classification data and mask information.</p> <ul style="list-style-type: none"> <li>• <b>port-string</b> - (Optional) Applies this administratively-assigned rule to a specific ingress port. 7100-Series devices with firmware versions 3.00.xx and higher also support the <b>set policy port</b> command as an alternative to administratively assign a profile rule to a port.</li> <li>• <b>storage-type</b> - (Optional) Adds or removes this entry from non-volatile storage.</li> <li>• <b>admin-pid</b> - Associates this administrative profile with a policy profile index ID. Valid values are 1 - 1023.</li> </ul>	<b>set policy rule admin-profile</b> { <b>macsource</b>   <b>port</b> } [ <i>data</i> ] [ <b>mask</b> <i>mask</i> ] <b>port-string</b> <i>port-string</i> [ <b>storage-type</b> { <b>non-volatile</b>   <b>volatile</b> }] [ <b>admin-pid</b> <i>admin-pid</i> ]

**Procedure 22-2 Configuring Classification Rules (continued)**

Step	Task	Command(s)
2.	<p>In any command mode, optionally configure policy rules to associate with a policy role.</p> <p>See <a href="#">Table 22-1</a> on page 22-7 for traffic classification-type descriptions.</p> <p>See the <b>set policy rule</b> command discussion in the command reference guide that comes with your device for traffic classification data and mask information.</p> <ul style="list-style-type: none"> <li>• <b>port-string</b> - (Optional) Applies this policy rule to a specific ingress port. 7100-Series devices also support the <b>set policy port</b> command as an alternative way to assign a profile rule to a port.</li> <li>• <b>storage-type</b> - (Optional) Adds or removes this entry from non-volatile storage.</li> <li>• <b>vlan</b> - (Optional) Classifies this rule to a VLAN ID.</li> <li>• <b>drop   forward</b> - (Optional) Specifies that packets within this classification will be dropped or forwarded.</li> <li>• <b>cos</b> - (Optional) Specifies that this rule will classify to a Class-of-Service ID. Valid values are 0 - 255. A value of -1 indicates that no CoS forwarding behavior modification is desired.</li> </ul>	<pre><b>set policy rule</b> <i>profile-index</i> <i>classification-type</i> [<i>data</i>] [<b>mask</b> <i>mask</i>] [<b>port-string</b> <i>port-string</i>] [<b>storage-type</b> {<b>non-volatile</b>   <b>volatile</b>}] [<b>vlan</b> <i>vlan</i>]   [<b>drop</b>   <b>forward</b>] [<b>admin-pid</b> <i>admin-pid</i>] [<b>cos</b> <i>cos</i>] [<b>quarantine-profile</b> <i>quarantine-profile</i>] [<b>clear-quarantine-profile</b>] [<b>prohibit-quarantine-profile</b>]</pre>
3.	(Optional) Change the system resource allocation policy profile from <b>default</b> to <b>router1</b> .	<pre><b>set system resource-alloc-policy</b> {<b>default</b>   <b>router1</b>}</pre>
4.	(Optional) Assigns a policy role to a port.	<pre><b>set policy port</b> <i>port-name</i> <i>admin-id</i></pre>

[Table 22-5](#) describes how to display policy information and statistics.

**Table 22-5 Displaying Policy Configuration and Statistics**

Task	Command(s)
In any command mode, display policy role information.	<pre><b>show policy profile</b> {<b>all</b>   <i>profile-index</i> [<i>consecutive-pids</i>] [<b>-verbose</b>]}</pre>
In any command mode, display the action the device should take if asked to apply an invalid or unknown policy, or the number of times the device has detected an invalid/unknown policy, or both action and count information.	<pre><b>show policy invalid</b> {<b>all</b>   <b>action</b>   <b>count</b>}</pre>
In any command mode, display VLAN-ID to policy role mappings table.	<pre><b>show policy mactable</b> <i>vlan-list</i></pre>
In any command mode, display policy classification and admin rule information.	<pre><b>show policy rule</b> <i>classification-type</i> [<i>data</i>] [<b>mask</b> <i>mask</i>] [<b>port-string</b> <i>port-string</i>] [<b>storage-type</b> {<b>non-volatile</b>   <b>volatile</b>}]   [<b>drop</b>   <b>forward</b>] [<b>dynamic-pid</b> <i>dynamic-pid</i>] [<b>cos</b> <i>cos</i>] [<b>admin-pid</b> <i>admin-pid</i>] [<b>-verbose</b>] [<b>-wide</b>]</pre>
In any command mode, display all policy classification capabilities for this device.	<pre><b>show policy capability</b></pre>


**Table 22-5 Displaying Policy Configuration and Statistics (continued)**

Task	Command(s)
In any command mode, display a list of currently supported traffic rules applied to the administrative profile for one or more ports.	<b>show policy allowed-type</b> <i>port-string</i> [-verbose]
In any command mode, display status of dynamically assigned roles.	<b>show policy dynamic</b> [override]

# Policy Configuration Example

This section presents a college-based policy configuration example. Figure 22-1 displays an overview of the policy configuration, with the 7100-Series switch in the data center server switch role. This overview display is followed by a complete discussion of the configuration example.

Figure 22-1 College-Based Policy Configuration





**Note:** This policy configuration example places the 7100-Series device within an overall policy context using the default system resource allocation policy profile. See the server[iSCSI] policy discussion for policy specific to the 7100-Series. For purposes of this discussion, Edge Switch and Distribution Switch refer to S-Series platforms, and the Data Center Server Switch refers to a 7100-Series switch.

## Roles

The example defines the following roles:

- **guest** – Used as the default policy for all unauthenticated ports. Connects a PC to the network providing internet only access to the network. Provides guest access to a limited number of the edge switch ports to be used specifically for internet only access. Policy is applied using the port level default configuration, or by authentication, in the case of the Services Edge Switch port internet only access PCs.
- **student** – Connects a dorm room PC to the network through a “Student” Fixed Switch port. A configured CoS rate limits the PC. Configured rules deny access to administrative and faculty servers. The PC authenticates using RADIUS. Hybrid authentication is enabled. The **student** policy role is applied using the filter-ID attribute. The base VLAN is applied using the tunnel attributes returned in the RADIUS response message. If all rules are missed, the settings configured in the **student** policy profile are applied.
- **phoneFS** – Connects a dorm room or faculty office VoIP phone to the network using a stackable fixed switch port. A configured CoS rate limits the phone and applies a high priority. The phone authenticates using RADIUS. Hybrid authentication is enabled. Policy is applied using the filter-ID returned in the RADIUS response message. The base VLAN is applied using the tunnel attributes returned in the RADIUS response message. If all rules are missed, the settings configured in the **phoneFS** policy profile are applied.
- **faculty** – Connects a faculty office PC to the network through a “Faculty” Fixed Switch port. A configured CoS rate limits the PC. A configured rule denies access to the administrative servers. The PC authenticates using RADIUS. Hybrid authentication is enabled. The **faculty** policy role is applied using the filter-ID attribute. The base VLAN is applied using the tunnel attributes returned in the RADIUS response message for the authenticating user. If all rules are missed, the settings configured in the **faculty** policy profile are applied.
- **phoneES** – Connects a services VoIP phone to the network using a Services Edge Switch port. A configured CoS rate limits the phone for both setup and payload, and applies a high priority. The phone authenticates using RADIUS. Tunnel authentication is enabled. The base VLAN is applied using the tunnel attributes returned in the RADIUS response message. Policy is applied using a mappable configuration. If all rules are missed, the settings configured in the **phoneES** policy profile are applied.
- **services** – Connects a services PC to the network through the Services Edge Switch port. A configured CoS rate limits the PC. Services are denied access to both the student and faculty servers. The PC authenticates using RADIUS. The base VLAN is applied using the tunnel attributes returned in the RADIUS response message for the authenticating user. The **services** policy role is applied using a policy mappable setting. The policy accounting, syslog, invalid action and TCI overwrite are enabled for this role. If all rules are missed, the settings configured in the **services** policy profile are applied.
- **distribution** – The Distribution policy role is applied at the Distribution Switch providing rate limiting.
- **server[iSCSI]** – The Server iSCSI policy role is applied to the Data Center Server Switch which provides low latency, high speed switching between the storage servers and the other servers in the data center.

## Policy Domains

It is useful to break up policy implementation into logical domains for ease of understanding and configuration. For this example, it is useful to consider five domains: basic edge, standard edge on the Fixed Switch, premium edge on the Services Edge Switch, premium distribution on the Distribution Switch, and data center on the Data Center Server Switch.

### Basic Edge

Protocols not appropriate to the edge should be blocked. For this example we will block DHCP, DNS, SNMP, SSH, Telnet and FTP at the edge on the data VLAN. We will forward destination port DHCP and DNS and source port for IP address request to facilitate auto configuration and IP address assignment. See [“Blocking Non-Edge Protocols at the Edge Network Layer”](#) on page 22-8 for a listing of protocols you should consider blocking at the edge.

### Standard Edge

Edge Switch platforms will be rate-limited using a configured CoS that will be applied to the student and faculty, and phoneFS policy roles. Fixed Switch support for hybrid authentication depends upon the platform and firmware release. The Fixed Switch in this example supports the hybrid authentication capability. Hybrid authentication will be enabled.

### Premium Edge

The Edge Switch will be rate-limited using a configured CoS that is applied to the services and phoneES policy role. This premium edge platform will be enabled for the following capabilities:

- Policy Accounting
- Syslog rule usage enabled and set to machine-readable
- Invalid policy action set to drop
- TCI overwrite enabled

### Premium Distribution

The Distribution Switch Router will be rate-limited using a configured CoS. Premium distribution will be enabled for the following policy capabilities:

- Policy Accounting
- Syslog Rule Usage enabled and set to machine-readable
- Invalid policy action set to drop
- TCI overwrite enabled

### Data Center

The Data Center Server Switch will provide policy for ports connecting iSCSI storage nodes with the other data center servers. This policy will allow forwarding of all TCP traffic on the iSCSI port 3260 with a CoS that provides low latency and high speed. It will also provide a bilateral set of rules that allow administrators to SSH to the switch on TCP port 22 and a destination rule to allow the node to SSH to another device.

## Platform Configuration

This section will provide the CLI based policy configuration on the following platforms:

- Student Fixed Switch
- Faculty Fixed Switch
- Services Edge Switch
- Distribution Switch
- Data Center Server Switch

In CLI mode, configuration takes place on each platform. When using the NetSight Policy Manager, configuration takes place at a central location and is pushed out to the appropriate network devices.

For this configuration example, CoS related configuration will be specified as a final CoS. For details on configuring CoS, see *“Understanding QoS Configuration on the 7100-Series”* on page 40-8.



**Note:** CLI command prompts used in this configuration example have the following meaning:

- Extreme Networks(rw)-> – Input on all platforms used in this example.
- Fixed Switch(rw)-> – Input on all Fixed Switches.
- StudentFS-> – Input on the student Fixed Switch.
- FacultyFS-> – Input on the faculty Fixed Switch.
- Services(rw)-> – Input on the services S-Series device.
- Distribution(rw)-> – Input on the distribution S-Series device.
- iSCSI(wr)-> – Input on the data center 7100-Series device.

## Configuring Guest Policy on Edge Platforms

All edge ports will be set with a default **guest** policy using the **set policy port** command. This guest policy provides for an internet only access to the network. Users on all ports will attempt to authenticate. If the authentication succeeds, the policy returned by authentication or, in the case of the Services Edge Switch configuration, the mappable setting, overrides the default port policy setting. If authentication fails, the guest policy is used. On the Services Edge Switch, five ports are used by PCs at locations throughout the campus, such as the library, to provide access to the internet. The PCs attached to these five ports will authenticate with the **guest** policy role. Public facing services would also be configured for guest status in a school or enterprise scenario. Public facing services are not part of this example.

### Configuring the Policy Role

The guest role is configured with:

- A profile-index value of **1**
- A name of **guest**
- A PVID set to **0**
- A CoS set to **4**

Create the guest policy profile on all platforms:

```
Extreme Networks(rw)->set policy profile 1 name guest pvid-status enable pvid 0
cos-status enable cos 4
```

## Assigning Traffic Classification Rules

For cases where discovery must take place to assign an IP address, DNS and DHCP traffic must be allowed. Forwarding of traffic is allowed on UDP source port 68 (IP address request) and UDP destination ports 53 (DNS) and 67 (DHCP).

```
Extreme Networks(rw)->set policy rule 1 udpsourceport 68 mask 16 forward
Extreme Networks(rw)->set policy rule 1 udpdestportIP 53 mask 16 forward
Extreme Networks(rw)->set policy rule 1 udpdestportIP 67 mask 16 forward
```

Guest policy allows internet traffic. TCP destination Ports 80, 8080, and 443 will be allowed traffic forwarding.

```
Extreme Networks(rw)->set policy rule 1 tcpdestportIP 80 mask 16 forward
Extreme Networks(rw)->set policy rule 1 tcpdestportIP 443 mask 16 forward
Extreme Networks(rw)->set policy rule 1 tcpdestport 8080 mask 16 forward
```

ARP forwarding is required on ether port 0x806.

```
Extreme Networks(rw)->set policy rule 1 ether 0x806 mask 16 forward
```

## Assigning the Guest Policy Profile to All Edge Ports

Assign the guest policy profile to all Fixed Switch and Services Edge Switch ports.

```
Extreme Networks(rw)->set policy port ge.*.1-47 1
```

## Configuring Policy for the Edge Student Fixed Switch

### Configuring the Policy Role

The student role is configured with:

- A profile-index value of **2**
- A name of **student**
- A port VLAN of **10**
- A CoS of **8**

Create a policy role that applies a CoS 8 to data VLAN 10 and configures it to rate-limit traffic to 1M with a moderate priority of 5.

```
StudentFS(rw)->set policy profile 2 name student pvid-status enable pvid 10
cos-status enable cos 8
```

### Assigning Hybrid Authentication

Configure the RADIUS server user accounts with the appropriate tunnel information using VLAN authorization and policy filter-ID for student role members and devices. Enable hybrid authentication, allowing the switch to use both the filter-ID and tunnel attributes in the RADIUS response message. Set a VLAN-to-policy mapping as backup in case the response does not include the RADIUS filter-ID attribute. This mapping is ignored if RADIUS filter-ID attribute is present in the RADIUS response message.

```
StudentFS(rw)->set policy mactable response both
StudentFS(rw)->set policy mactable 10 2
```

### Assigning Traffic Classification Rules

Forward traffic on UDP source port for IP address request (68), and UDP destination ports for protocols DHCP (67) and DNS (53). Drop traffic on UDP source ports for protocols DHCP (67) and DNS (53). Drop traffic for protocols SNMP (161), SSH (22), Telnet (23) and FTP (20 and 21) on both the data and phone VLANs.



```
StudentFS(rw)->set policy rule 2 udpsourceport 68 mask 16 forward
StudentFS(rw)->set policy rule 2 udpdestport 67 mask 16 forward
StudentFS(rw)->set policy rule 2 udpdestport 53 mask 16 forward
StudentFS(rw)->set policy rule 2 udpsourceportIP 67 mask 16 drop
StudentFS(rw)->set policy rule 2 udpsourceportIP 53 mask 16 drop
StudentFS(rw)->set policy rule 2 udpdestportIP 16 mask 16 drop
StudentFS(rw)->set policy rule 2 tcpdestportIP 22 mask 16 drop
StudentFS(rw)->set policy rule 2 tcpdestportIP 23 mask 16 drop
StudentFS(rw)->set policy rule 2 tcpdestportIP 20 mask 16 drop
StudentFS(rw)->set policy rule 2 tcpdestportIP 21 mask 16 drop
```

Students should only be allowed access to the services server (subnet 10.10.50.0/24) and should be denied access to both the administrative (subnet 10.10.60.0/24) and faculty servers (subnet 10.10.70.0/24).

```
StudentFS(rw)->set policy rule 2 ipdestsocket 10.10.60.0 mask 24 drop
StudentFS(rw)->set policy rule 2 ipdestsocket 10.10.70.0 mask 24 drop
```

## Configuring PhoneFS Policy for the Edge Fixed Switch

### Configuring the Policy Role

The phoneFS role is configured on both the dorm room and faculty office Fixed Switches with:

- A profile-index of 3
- A name of **phoneFS**
- A port VLAN of 11
- A CoS of 10

Because we can not apply separate rate limits to the phone setup and payload ports on the Fixed Switch using policy rules, apply CoS 10 with the higher payload appropriate rate limit of 100k bps and a high priority of 6 to the phoneFS role.

```
Fixed Switch(rw)->set policy profile 3 name phoneFS pvid-status enable pvid 11
cos-status enable cos 10
```

### Assigning Traffic Classification Rules

Drop traffic for protocols SNMP (161), SSH (22), Telnet (23) and FTP (20 and 21) on the phone VLAN. Forward traffic on UDP source port for IP address request (68) and forward traffic on UDP destination ports for protocols DHCP (67) and DNS (53) on the phone VLAN, to facilitate phone auto configuration and IP address assignment.

```
Fixed Switch(rw)->set policy rule 3 udpdestportIP 161 mask 16 drop
Fixed Switch(rw)->set policy rule 3 tcpdestportIP 22 mask 16 drop
Fixed Switch(rw)->set policy rule 3 tcpdestportIP 23 mask 16 drop
Fixed Switch(rw)->set policy rule 3 tcpdestportIP 20 mask 16 drop
Fixed Switch(rw)->set policy rule 3 tcpdestportIP 21 mask 16 drop
Fixed Switch(rw)->set policy rule 3 udpsourceport 68 mask 16 forward
Fixed Switch(rw)->set policy rule 3 udpdestportIP 67 mask 16 forward
Fixed Switch(rw)->set policy rule 3 udpdestportIP 53 mask 16 forward
```

### Assigning Hybrid Authentication

Configure the RADIUS server user accounts with the appropriate tunnel information using VLAN authorization and policy filter-ID for phoneFS role members and devices. Enable hybrid authentication, allowing the switch to use both the filter-ID and tunnel attributes in the RADIUS response message. Set a VLAN-to-policy mapping as backup in case the response does not include the RADIUS filter-ID attribute. This mapping is ignored if RADIUS filter-ID attribute is present in the RADIUS response message.

```
Fixed Switch(rw)->set policy mactable response both
Fixed Switch(rw)->set policy mactable 11 3
```

## Configuring Policy for the Edge Faculty Fixed Switch

### Configuring the Policy Role

The faculty role is configured with:

- A profile-index value of **4**
- A name of **faculty**
- A port VLAN of **10**
- A CoS of **8**

Create a policy role that applies a CoS 8 to data VLAN 10 and configures it to rate-limit traffic to 1M with a moderate priority of 5.

```
FacultyFS(rw)->set policy profile 4 name faculty pvid-status enable pvid 10
cos-status enable cos 8
```

### Assigning Hybrid Authentication

Configure the RADIUS server user accounts with the appropriate tunnel information using VLAN authorization and policy filter-ID for faculty role members and devices. Enable hybrid authentication. Set a VLAN-to-policy mapping. This mapping is ignored if the RADIUS filter-ID attribute is present in the RADIUS response message.

```
StudentFS(rw)->set policy mactable response both
StudentFS(rw)->set policy mactable 10 4
```

### Assigning Traffic Classification Rules

Forward traffic on UDP source port for IP address request (68), and UDP destination ports for protocols DHCP (67) and DNS (53). Drop traffic on UDP source ports for protocols DHCP (67) and DNS (53). Drop traffic for protocols SNMP (161), SSH (22), Telnet (23) and FTP (20 and 21) on both the data and phone VLANs.

```
FacultyFS(rw)->set policy rule 4 udpsourceport 68 mask 16 forward
FacultyFS(rw)->set policy rule 4 udpdestport 67 mask 16 forward
FacultyFS(rw)->set policy rule 4 udpdestport 53 mask 16 forward
FacultyFS(rw)->set policy rule 4 udpsourceportIP 67 mask 16 drop
FacultyFS(rw)->set policy rule 4 udpsourceportIP 53 mask 16 drop
FacultyFS(rw)->set policy rule 4 udpdestportIP 16 mask 16 drop
FacultyFS(rw)->set policy rule 4 tcpdestportIP 22 mask 16 drop
FacultyFS(rw)->set policy rule 4 tcpdestportIP 23 mask 16 drop
FacultyFS(rw)->set policy rule 4 tcpdestportIP 20 mask 16 drop
FacultyFS(rw)->set policy rule 4 tcpdestportIP 21 mask 16 drop
```

Faculty should only be allowed access to the services (subnet 10.10.50.0/24) and the faculty servers (subnet 10.10.70.0/24) and should be denied access to the administrative server (subnet 10.10.60.0/24).

```
FacultyFS(rw)->set policy rule 4 ipdestsocket 10.10.60.0 mask 24 drop
```

## Configuring PhoneES Policy for the Services Edge Switch

### Configuring the Policy Role

The phoneES role is configured on the Services Edge Switch with:

- A profile-index of **5**

- A name of **phoneES**
- A default port VLAN of **0**
- A default CoS of **4**

Because VLANs can be applied to Services Edge Switch ports using the appropriate traffic classification, the explicit deny all PVID **0** will be applied at policy creation. Separate rate limits can be applied to the phone setup and payload ports on the Services Edge Switch using policy rules. A default CoS of 4 will be applied at policy role creation.

```
ServicesES(rw)->set policy profile 5 name phoneES pvid-status enable pvid 0
cos-status enable cos 4
```

### Assigning Traffic Classification Rules

Forward traffic on UDP source port for IP address request (68) and forward traffic on UDP destination ports for protocols DHCP (67) and DNS (53) on the phone VLAN, to facilitate phone auto configuration and IP address assignment. Drop traffic for protocols SNMP (161), SSH (22), Telnet (23) and FTP (20 and 21) on the phone VLAN.

```
ServicesES(rw)->set policy rule 5 udpsourceport 68 mask 16 forward
ServicesES(rw)->set policy rule 5 udpdestportIP 67 mask 16 forward
ServicesES(rw)->set policy rule 5 udpdestportIP 53 mask 16 forward
ServicesES(rw)->set policy rule 5 udpdestportIP 161 mask 16 drop
ServicesES(rw)->set policy rule 5 tcpdestportIP 22 mask 16 drop
ServicesES(rw)->set policy rule 5 tcpdestportIP 23 mask 16 drop
ServicesES(rw)->set policy rule 5 tcpdestportIP 20 mask 16 drop
ServicesES(rw)->set policy rule 5 tcpdestportIP 21 mask 16 drop
```

Apply a CoS 9 to phone setup data on VLAN 11, rate limiting the data to 5 pps with a high priority of 7 on port 2427.

Apply a CoS 10 to phone payload data on VLAN 11, rate limiting the data to 100k bps with a high priority of 7 for both source and destination on port 5004.

```
ServicesES(rw)->set policy rule 5 upddestIP 2427 mask 16 vlan 11 cos 9
ServicesES(rw)->set policy rule 5 udpsourceIP 5004 mask 16 vlan 11 cos 10
ServicesES(rw)->set policy rule 5 upddestIP 5004 mask 16 vlan 11 cos 10
```

### Assigning the VLAN-to-Policy Association

The nature of services related devices that might connect to a switch port is not as static as with the student or faculty roles. Services related network needs can run the gamut from temporary multimedia events to standard office users. There may be multiple VLAN and policy role associations that take care of services related needs, depending upon the connected user. This may include the requirement for multiple services related roles.

For services, the network administrator desires greater resource usage flexibility in assigning the policy to VLAN association. Authentication in this case will return only the tunnel attributes in the response message based upon the requirements of the authenticating user. Setting the VLAN-to-policy association will be handled by the mactable configuration, allowing for ease in changing the policy associated with a VLAN on the fly using Policy Manager. Specify that the **tunnel** attributes returned in the RADIUS response message will be used by the authenticating user. Associate VLAN 11 with policy role 5 using the **set policy mactable** command.

```
ServicesES(rw)->set policy mactable response tunnel
ServicesES(rw)->set policy mactable 11 5
```

## Configuring Policy for the Services Edge Switch

### Configuring the Policy Role

The services role is configured with:

- A profile-index value of 6
- A name of **services**
- A default port VLAN of 0
- A default CoS when no rule overrides CoS
- TCI overwrite enabled

```
ServicesES(rw)->set policy profile 6 name services pvid-status enable pvid 0
cos-status enable cos 4 tci-overwrite enable
```

### Assigning the VLAN-to-Policy Association

Setting the VLAN-to-policy association will be handled by the policy mactable setting, allowing for ease in changing the policy associated with a VLAN on the fly using Policy Manager. Specify that the **tunnel** attributes returned in the RADIUS response message will be used by the authenticating user. Associate VLAN 10 with policy role 6 using the **set policy mactable** command.

```
ServicesES(rw)->set policy mactable response tunnel
ServicesES(rw)->set policy mactable 10 6
```

### Assigning Traffic Classification Rules

Forward traffic on UDP source port for IP address request (68) and forward traffic on UDP destination ports for protocols DHCP (67) and DNS (53) on the data VLAN, to facilitate PC auto configuration and IP address assignment. Drop traffic for protocols SNMP (161), SSH (22), Telnet (23) and FTP (20 and 21) on the phone VLAN.

```
ServicesES(rw)->set policy rule 6 udpsourceportIP 68 mask 16 vlan 10 forward
ServicesES(rw)->set policy rule 6 udpdestportIP 67 mask 16 vlan 10 forward
ServicesES(rw)->set policy rule 6 udpdestportIP 53 mask 16 vlan 10 forward
ServicesES(rw)->set policy rule 6 udpdestportIP 67 mask 16 vlan 10 drop
ServicesES(rw)->set policy rule 6 udpdestportIP 53 mask 16 vlan 10 drop
ServicesES(rw)->set policy rule 6 udpdestportIP 161 mask 16 drop
ServicesES(rw)->set policy rule 6 tcpdestportIP 22 mask 16 drop
ServicesES(rw)->set policy rule 6 tcpdestportIP 23 mask 16 drop
ServicesES(rw)->set policy rule 6 tcpdestportIP 20 mask 16 drop
ServicesES(rw)->set policy rule 6 tcpdestportIP 21 mask 16 drop
```

Apply a CoS 8 to data VLAN 10 and configure it to rate-limit traffic to 1M and moderate priority of 5 for services IP subnet 10.10.30.0 mask 28. We will also enable traps and syslog for this subnet.

```
ServicesES(rw)->set policy rule 6 ipsourcesocket 10.10.30.0 mask 28 syslog enable
trap enable vlan 10 cos 8
```

Services should only be allowed access to the services server (subnet 10.10.50.0/24) and should be denied access to the faculty servers (subnet 10.10.70.0/24) and administrative servers (subnet 10.10.60.0/24).

```
ServicesES(rw)->set policy rule 6 ipdestsocket 10.10.60.0 mask 24 drop
ServicesES(rw)->set policy rule 6 ipdestsocket 10.10.70.0 mask 24 drop
```

### Enable Enhanced Edge Switch Capabilities on the Services Edge Switch Platform

The Services Edge Switch platform supports a number of enhanced capabilities not available on the Fixed Switch platforms. The following enhanced policy capabilities are enabled: policy

accounting to flag the occurrence of a rule hit, syslog rule usage set to machine-readable for enterprise specific backend syslog statistics gathering, an invalid action set to default policy should an invalid policy occur, TCI overwrite enabled to allow for Type of Service (ToS) overwrite for the VoIP phone.

```
ServicesES(rw)->set policy accounting enable
ServicesES(rw)->set policy syslog machine-readable
ServicesES(rw)->set policy invalid action default-policy
ServicesES(rw)->set port tcioverwrite ge.1.1-10
```

## Configuring the Distribution Layer Role

### Configuring the Policy Role

The distribution role is configured with:

- A profile-index value of 7
- A name of **distribution**
- A default CoS when no rule overrides CoS
- TCI overwrite enabled

```
Distribution(rw)->set policy profile 7 name distribution cos-status enable cos 4
tcioverwrite enable
```

### Assigning the Traffic Classification to the Policy Role

Assign ports ge.1.1-26 to the distribution policy role, specifying the associated ports **1 - 26**, enable traps and enable syslog.

```
Distribution(rw)->set policy rule admin-profile port ge.1.1-26 admin-pid 7
port-string ge.1.1-26 trap enable syslog enable.
```

### Assigning Traffic Classification Rules

Assign a CoS to distribution up and down stream link ports, rate-limiting the traffic to 25M.

```
Distribution(rw)->set policy rule 7 port ge.1.1-26 cos 11
Distribution(rw)->set policy rule 7 port ge.2.1-26 cos 11
```

### Enable Enhanced Policy Capabilities on the Distribution Platform

The Distribution platform supports a number of policy capabilities not available on the Fixed Switch platforms. The following enhanced policy capabilities are enabled: policy accounting to flag the occurrence of a rule hit, syslog rule usage set to machine-readable for backend syslog statistics gathering, an invalid action set to default policy should an invalid policy occur, TCI overwrite enabled to allow for Type of Service (ToS) overwrite for the VoIP phone.

```
ServicesES(rw)->set policy accounting enable
ServicesES(rw)->set policy syslog machine-readable
ServicesES(rw)->set policy invalid action default-policy
ServicesES(rw)->set port tcioverwrite ge.1.1-26
ServicesES(rw)->set port tcioverwrite ge.2.1-26
```

## Configuring Server[iSCSI] Policy on the 7100-Series Platform

Servers will access iSCSI storage by communicating with iSCSI storage nodes in the server farm through 7100-Series ports configured with the server[iSCSI] policy role. This policy will allow forwarding of all TCP traffic on the iSCSI port 3260 with a CoS that provides low latency and high speed. It will also provide a bilateral set of rules that allow administrators to SSH to the switch on TCP port 22 and a destination rule to allow the node to SSH to another device.

## Configuring the Policy Role

The `server[iSCSI]` role is configured with:

- A profile-index of `12`
- A name of `server[iSCSI]`
- Ports `tg.1.10-15`
- PVID `0`
- CoS `12`

Create the `server[iSCSI]` role on the Data Center Server Switch with a default action of deny all (PVID `0`):

```
iSCSI(rw)->set policy profile 12 name "server[iSCSI]" pvid-status enable pvid 0
```

## Assigning Traffic Classification Rules

Allow the server farm storage nodes to communicate on TCP source port `3260` with a CoS `12` that prioritizes the traffic for low latency and high speed.

```
iSCSI(rw)->set policy rule 1 tcpsourceportIP 3260 mask 16 forward cos 12
```

Allow administrator access to the device using SSH on TCP source port `22` and the node to SSH to another device on TCP destination port `22`.

```
set policy rule 1 tcpsourceportIP 22 mask 16 forward
set policy rule 1 tcpsourceportIP 22 mask 16 forward
```

Apply this profile to ports `tg.1.10-15`.

```
set policy rule admin-profile port tg.1.10-15 mask 16 port-string tg.1.10-15
admin-pid 1
```

This completes the policy configuration for this school example.

## Terms and Definitions

[Table 22-6](#) lists terms and definitions used in this policy configuration discussion.

**Table 22-6 Policy Configuration Terms and Definitions**

Term	Definition
Administrative Profile	A logical container that assigns a traffic classification to a policy role.
Class of Service (CoS)	A logical container for packet priority, queue, and forwarding treatment that determines how the firmware treats a packet as it transits the link.
Filter-ID	A string that is formatted in the RADIUS access-accept packet sent back from the authentication server to the switch during the authentication process. In the Extreme Networks policy context, the string contains the name of the policy role to be applied to the authenticating user or device.
Hybrid Authentication	An authentication feature that allows the switch to use both the filter-ID and tunnel attributes in the RADIUS response message to determine how to treat the authenticating user.
Policy	A component of Secure Networks that provides for the configuration of a role based profile for the securing and provisioning of network resources based upon the function the user or device plays within the enterprise network.
Policy Mappable	A logical entity that can be configured to provide VLAN to policy role mappings.

**Table 22-6 Policy Configuration Terms and Definitions (continued)**

<b>Term</b>	<b>Definition</b>
Policy Profile/Role	A logical container for the rules that define a particular policy role.
Policy Rule	A logical container providing for the specification of policy behaviors associated with a policy role.
Role	The grouping of individual users or devices into a logical behavioral profile for the purpose of applying policy.
System Resource Allocation Policy	A set of profiles that determine the traffic classifications and ACL types supported for the system and the number of rules supported for each traffic classification and ACL type
Traffic Classification	A network element such as MAC or IP address, packet type, port, or VLAN used as the basis for identifying the traffic to which the policy will be applied.
Untagged and Tagged VLAN	Untagged VLAN frames are classified to the VLAN associated with the port it enters. Tagged VLAN frames are classified to the VLAN specified in the VLAN tag; the PVID is ignored.
VLAN Authorization	An aspect of RFC3580 that provides for the inclusion of the VLAN tunnel attribute in the RADIUS Access-Accept packet defining the base VLAN-ID to be applied to the authenticating user or device.
VLAN Egress List	A configured list of ports that a frame for this VLAN can exit.

## Multicast Configuration

This document describes the multicast feature and its configuration on 7100-Series devices.

For information about...	Refer to page...
<a href="#">How to Use Multicast in Your Network</a>	23-1
<a href="#">Implementing Multicast</a>	23-2
<a href="#">Understanding Multicast</a>	23-2
<a href="#">Configuring Multicast</a>	23-20

### How to Use Multicast in Your Network

Multicast is a “one source to many destinations” method of simultaneously sending information over a network using the most efficient delivery strategy over each link. Only the end stations that explicitly indicate a need to receive a given multicast stream will receive it.

Applications that take advantage of multicast include video conferencing, streaming video, corporate communications, distance learning, and distribution of software, stock quotes, and news.

Multicast technology includes the following protocols:

- Internet Group Management Protocol (IGMP) for IPv4, Multicast Listener Discovery (MLD) for IPv6
- Distance Vector Multicast Routing Protocol (DVMRP)
- Protocol Independent Multicast (PIM)

Unlike unicast and broadcast, multicast uses network infrastructure efficiently because only one copy of the source traffic is sent throughout the network, going only to interested receivers, minimizing the burden placed on the sender, network, and receiver. The routers in the network take care of replicating the packet, where necessary, to reach multiple receivers. If a router decides that there are no interested users downstream from itself, it prunes the stream back to the next router. Thus, unwanted streams are not sent to the pruned routers, saving bandwidth and preventing unwanted packets from being sent.



## Implementing Multicast

You can implement the IGMP, DVMRP, and PIM multicast protocols on Extreme Networks devices using simple CLI commands as described in this document. A basic configuration process involves the following tasks:

1. Configuring the VLANs and IP interfaces on which you want to transmit multicast.
2. Enabling the multicast protocol(s) on configured interfaces.

For PIM, you must also configure a unicast routing protocol, such as OSPF. For both DVMRP and PIM for IPv4 to operate, IGMP must be enabled. For PIM for IPv6 to operate, the Multicast Listener Discovery (MLD) protocol must be enabled.

## Understanding Multicast

Multicast allows a source to send a single copy of data using a single IP address from a well-defined range for an entire group of recipients (a multicast group). A source sends data to a multicast group by simply setting the destination IP address of the datagram to be the multicast group address. Sources do not need to register in any way before they can begin sending data to a group, and do not need to be members of the group themselves. Routers between the source and recipients use the group address to route the data, forwarding duplicate data packets only when the path to recipients diverges.

Hosts that wish to receive data from the multicast group join the group by sending a message to a multicast router on a local interface, using a multicast group membership discovery protocol, such as IGMP (IPv4) or MLD (IPv6). For more IGMP information, see [“Internet Group Management Protocol \(IGMP\)”](#) on page 23-2. For more MLD information, see [Chapter 24, Multicast Listener Discovery \(MLD\) Configuration](#).

Multicast routers communicate among themselves using a multicast routing protocol, such as DVMRP, PIM-SM, or PIM-DM. These protocols calculate a multicast distribution tree (PIM-SM) or source-based tree (DVMRP and PIM-DM) of recipients to ensure that:

- Multicast traffic reaches all recipients that have joined the multicast group
- Multicast traffic does not reach networks that do not have any such recipients (unless the network is a transit network on the way to other recipients)
- The number of identical copies of the same data flowing over the same link is minimized

For more information, see [“Distance Vector Multicast Routing Protocol \(DVMRP\)”](#) on page 23-5 and [“Protocol Independent Multicast \(PIM\)”](#) on page 23-11.

## Internet Group Management Protocol (IGMP)

### Overview

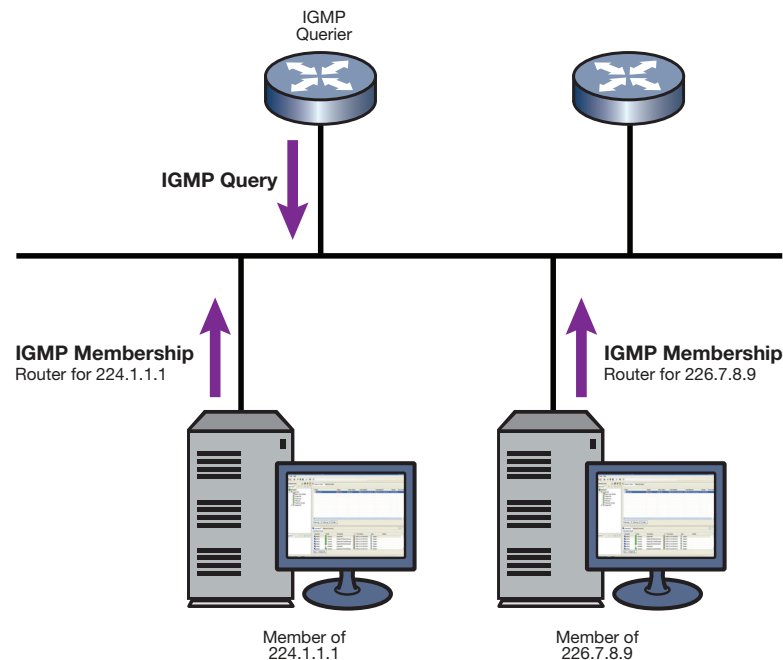
Group membership management is fundamental to the multicasting process. An arbitrary group of receivers can express interest in receiving a particular multicast stream, regardless of the physical or geographical boundaries of its members.

The purpose of IP multicast group management is to optimize a switched network’s performance so multicast packets will only be forwarded to those ports containing multicast group hosts or multicast switch devices instead of flooding to all ports in the subnet (VLAN).

IGMP uses three key components to control multicast membership:

- **Source** — A server that sends an IP multicast data stream with a particular multicast destination IP and MAC address. A server may not have direct IGMP involvement, as it often does not receive a multicast stream, but only sends a multicast stream.
- **Querier** — A device that periodically sends out queries in search of multicast hosts on a directly connected network. If multiple queriers are present on the LAN, the querier with the lowest IP address assumes the role.
- **Host** — A client end station that sends one of two IGMP messages to a querier:
  - Join message — Indicates the host wants to receive transmissions associated to a particular multicast group.
  - Leave message — Indicates the host wants to stop receiving the multicast transmissions.

**Figure 23-1 IGMP Querier Determining Group Membership**



As shown in [Figure 23-1](#), a multicast query-enabled device can periodically ask its hosts if they want to receive multicast traffic. If there is more than one device on the LAN performing IP multicasting, one of these devices is elected querier and assumes the responsibility of querying the LAN for group members.

Based on the group membership information learned from IGMP, a device can determine which (if any) multicast traffic needs to be forwarded to each of its ports. At Layer 3, multicast switch devices use this information, along with a multicast routing protocol, to support IP multicasting across the Internet.

IGMP provides the final step in IP multicast delivery. It is only concerned with forwarding multicast traffic from the local switch device to group members on a directly attached subnetwork or LAN segment.

IGMP neither alters nor routes any IP multicast packets. Since IGMP is not concerned with the delivery of IP multicast packets across subnetworks, an external IP multicast device is needed if IP multicast packets have to be routed across different subnetworks.



**Note:** On VLANs where IGMP snooping is enabled, any received multicast stream will be flooded to the VLAN until such time as the IGMP database is populated, then stream forwarding will revert to ports with group membership only.

## IGMP Support on Extreme Networks Devices

Extreme Networks devices implement IGMP version 2 (RFC 2236) and IGMP version 3 (RFC 3376), which includes interoperability with version 1 hosts. IGMP version 1 is defined in RFC 1112.

Depending on your Extreme Networks device, IGMP can be configured independently at the switch level (Layer 2) and at the router level (Layer 3).

Extreme Networks devices support IGMP as follows:

- Passively snooping on the IGMP query and IGMP report packets transferred between IP multicast switches and IP multicast host groups to learn IP multicast group members. Each Layer 2 device records which ports IGMP packets are received on, depending on the kind of IGMP message, so multicast data traffic is not flooded across every port on the VLAN when it is received by the switch.

IGMP querying is disabled by default, use **set igmp query-enable** command to enable IGMP querying as described in “Configuring IGMP” on page 23-20.

- Actively sending IGMP query messages to learn locations of multicast switches and member hosts in multicast groups within each VLAN.
- Configuration of static IGMP groups which provides for specifying the IP address (group address) and VLAN of a non-IGMP capable device, forcing the sending of multicast traffic to the device. You can configure a static IGMP group using the **set igmp static** command as described in “Configuring IGMP” on page 23-20.

### Example: Sending a Multicast Stream

Figure 23-2 Sending a Multicast Stream with No Directly Attached Hosts

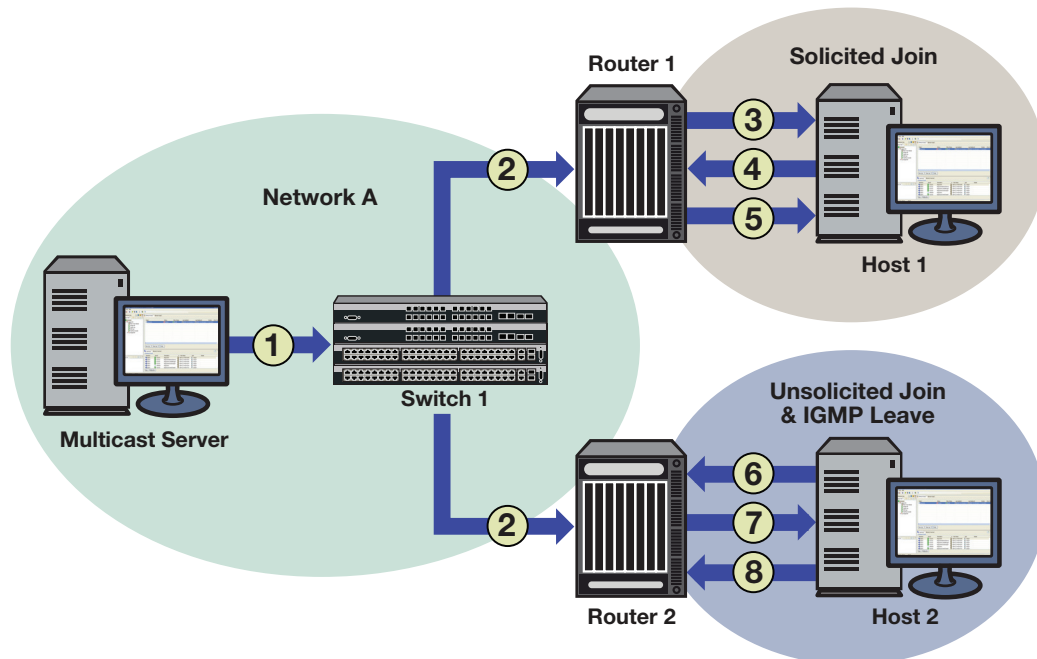


Figure 23-2 provides an example of IGMP processing on Extreme Networks devices when there are no directly attached hosts.

1. A single IP multicast server, with no directly attached hosts, sends a multicast stream into the network via Switch 1.

2. Because IGMP snooping is disabled, Switch 1 floods the multicast stream to all ports the multicast stream to all ports on the ingress VLAN.

Each router performs an IGMP forwarding check to see if there are any hosts that want to join the multicast group on its locally attached network. Each router drops multicast packets until a host joins the group using one of the following messages:

- **solicited join** (sent in response to an IGMP query produced by the router's interface)

In [Figure 23-2](#), this type of exchange occurs between Router 1 and Host 1 when:

- (3) Router 1 sends a query to potential Host 1.
- (4) Host 1 responds with a join message.
- (5) Router 1 forwards the multicast stream.

- **unsolicited join** (sent as a request without receiving an IGMP query first)

In [Figure 23-2](#), this type of exchange occurs between Router 2 and Host 2 when:

- (6) Host 2 sends a join message to Router 2.
- (7) Router 2 forwards the multicast stream to Host 2.
- (8) When it no longer wants to receive the stream, Host 2 can do one of the following:
  - Send a leave message to Router 2.
  - Time out the IGMP entry by not responding to further queries from Router 2.

## Distance Vector Multicast Routing Protocol (DVMRP)

### Overview

DVMRP, which is used for routing multicasts within a single, autonomous system, is designed to be used as an interior gateway protocol (IGP) within a multicast domain. It is a distance-vector routing protocol that relies on IGMP functionality to provide connectionless datagram delivery to a group of hosts across a network.



**Note:** IGMP must be enabled for DVMRP to operate.

DVMRP routes multicast traffic using a technique known as reverse path forwarding (RPF). When a router receives IP multicast packets, it first does an RPF check to determine if the packets are received on the correct interface. If so, the router forwards the packets out to the following:

- Local IGMP receivers for that group on interfaces for which the transmitting router is the designated forwarder
- Neighbor routers that have indicated their dependence on the transmitting router for forwarding multicast packets from that source (this is determined during DVMRP Route Exchange) and from which the transmitting router has not received any prune messages.

If not, the packets are discarded by the router. The transmitting router does not forward the packets back to the source.

If a router is attached to a set of VLANs that do not want to receive from a particular multicast group, the router can send a prune message back up the distribution tree to stop subsequent packets from traveling where there are no members. DVMRP periodically re-floods in order to reach any new hosts that want to receive from a particular group.

DVMRP routers dynamically discover their neighbors by sending neighbor probe messages periodically to an IP multicast group address that is reserved for all DVMRP routers.

Key features of DVMRP are the following:

- uses the well-known multicast IP address 224.0.0.4
- uses IGMP to exchange routing datagrams
- does not require an underlying Layer 3 routing protocol to provide a path to remote multicast destinations
- combines many of the features of the Routing Information Protocol (RIP) with the Truncated Reverse Path Broadcasting (TRPB) algorithm to route multicast packets between sources and receivers

## DVMRP Support on Extreme Networks Devices

DVMRP routing is implemented on Extreme Networks devices as specified in RFC 1075 and *draft-ietf-idmr-dvmrp-v3-10.txt*.

Extreme Networks devices support the following DVMRP components:

- [Probe Messages](#) for neighbor discovery
- [Route Table](#) for maintaining routes to all DVRMP networks
- [Route Reports](#) for route exchange with adjacent devices
- [Mroute Table](#) for maintaining per-source-group multicast trees
- [Prune Messages](#) for terminating multicast delivery trees
- [Graft Messages](#) for re-adding pruned multicast delivery trees

### Probe Messages

Each DVMRP-enabled interface transmits multicast probe packets to inform other DVMRP routers that it is operational. Probe messages are sent every 10 seconds on every interface running DVMRP. These messages provide:

- **A mechanism for DVMRP devices to locate each other.** Probe messages contain a list of the neighbors detected for each enabled interface. If no neighbors are found, the network is considered to be a leaf network.
- **A mechanism for DVMRP devices to determine the capabilities of neighboring devices.** Probe messages contain flags about neighbors' DVMRP capabilities and version compliance.
- **A keep-alive function for quickly detecting neighbor loss.** If a probe message from an adjacent neighbor is not seen within 35 seconds, the neighbor is timed out.

### Route Table

Each DVMRP-enabled device builds a DVMRP route table to maintain routes to all networks involved in DVMRP routing. As shown in the following 7100-Series modular switch example, the DVMRP route table contains a destination and next hop IP address, associated interface, metric value, expiration time, and up-time.

```
7100 Chassis(su)->show ip dvmrp route
```

Destination	Next Hop	Interface	Metric	Expire	Uptime
9.9.9.0/24	168.3.2.1	vlan.0.3200	3	00:01:52	2d, 19:34:45
21.2.2.0/24	168.3.2.1	vlan.0.3200	2	00:01:52	3d, 01:14:49
21.21.21.0/24	168.3.2.1	vlan.0.3200	2	00:01:52	3d, 01:14:49
29.2.2.0/24	168.3.2.1	vlan.0.3200	2	00:01:52	3d, 01:14:49

32.1.1.0/24	168.3.2.1	vlan.0.3200	2	00:01:52	3d, 01:14:49
32.11.11.0/24	168.3.2.1	vlan.0.3200	2	00:01:52	3d, 01:14:49
92.9.2.0/24	168.3.2.1	vlan.0.3200	2	00:01:52	3d, 01:14:49
100.3.3.0/24	Connected	vlan.0.3200	1	00:00:00	02:09:22
129.2.9.0/24	168.3.2.1	vlan.0.3200	2	00:01:52	2d, 19:02:06
139.3.9.0/28	Connected	vlan.0.390	1	00:00:00	3d, 01:14:54
160.2.2.0/24	168.3.2.1	vlan.0.3200	2	00:01:52	3d, 01:14:49
168.2.1.0/24	168.3.2.1	vlan.0.3200	2	00:01:52	3d, 01:14:49
168.3.0.0/16	Connected	vlan.0.3200	1	00:00:00	02:09:22
168.3.1.0/26	Connected	vlan.0.3100	5	00:00:00	2d, 21:54:44
168.8.1.0/24	168.3.2.1	vlan.0.3200	3	00:01:52	2d, 19:34:25
188.21.21.0/24	168.3.2.1	vlan.0.3200	2	00:01:52	3d, 01:14:49
188.23.23.0/24	168.3.2.1	vlan.0.3200	2	00:02:02	3d, 01:14:49
189.8.9.0/24	168.3.2.1	vlan.0.3200	4	00:02:02	2d, 19:34:15
191.9.1.0/24	168.3.2.1	vlan.0.3200	3	00:02:02	2d, 19:34:45
191.9.9.0/24	168.3.2.1	vlan.0.3200	3	00:02:02	2d, 19:34:45
192.9.2.0/24	168.3.2.1	vlan.0.3200	2	00:02:02	3d, 01:14:49
193.9.3.0/24	Connected	vlan.0.930	1	00:00:00	3d, 01:14:54
198.9.8.0/24	168.3.2.1	vlan.0.3200	4	00:02:02	2d, 19:34:15
198.23.23.0/24	168.3.2.1	vlan.0.3200	2	00:02:02	3d, 01:14:49
199.23.23.0/24	168.3.2.1	vlan.0.3200	2	00:02:02	3d, 01:14:49
250.9.9.0/24	168.3.2.1	vlan.0.3200	2	00:02:02	3d, 01:14:49

The number of DVMRP routes is 26

## Route Reports

DVMRP-enabled devices send route report packets to adjacent DVMRP devices every 60 seconds. When a DVMRP device receives one, it checks to verify that the report is from a known neighbor before processing.

The first time a device sees its own address in a neighbor's probe packet, it sends a unicast copy of its entire routing table to the neighbor to reduce start-up time.

The route report packet contains data about all networks/routes of which the sending device is aware. This information is used to determine the reverse path back to a particular multicast source. Every DVMRP device keeps a separate metric associated with each route. This metric is the sum of all interface metrics between the device originating the report and the source network.

DVMRP devices accept route reports for aggregated source networks in accordance with classless inter-domain devices (CIDR). This means that, if a prune or graft is received on a downstream interface for which the source network is aggregated, then a prune or graft should be sent upstream (to the multicast source).

If a DVMRP device has a large number of DVMRP routes, it will spread route reports across the route update interval (60 seconds) to avoid bottlenecks in processing and route synchronization issues.

For the purpose of pruning, DVMRP needs to know which downstream routes depend on the device for receiving multicast streams. Using poison reverse, the upstream router maintains a table of the source network and all downstream devices that are dependent on the upstream device.

## Mroute Table

DVMRP-enabled devices use the mroute table to maintain a source-specific forwarding tree.

When a DVMRP device is initialized, it assumes the role of the designated forwarder for all of its locally attached networks. Before forwarding any packets, all devices use IGMP to learn which networks would like to receive particular multicast group streams. In the case of a shared network, the device with a lower interface metric (a configurable value), or the lower IP address will become the designated forwarder.

A DVMRP device forwards multicast packets first by determining the upstream interface, and then by building the downstream interface list. If a downstream router has no hosts for a multicast stream, it sends a prune message to the upstream router. If the upstream router's outbound list is now empty, it may send a prune message to its upstream router.

If a downstream device has pruned a multicast group that a host would like to now receive, the downstream device must send a DVMRP graft message to its upstream device. The DVMRP graft will traverse the source-specific multicast delivery tree to the device that is receiving this stream.

As shown in the following example, the Mroute table displays the incoming interface IP address, the multicast group address, the uptime of the stream, incoming interface port number, and the outgoing interface port number.

```
7100 Chassis(su)->show ip mroute
```

```
IP Multicast Routing Table
```

```
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned
```

```
       R - RP-bit set, F - Register flag, T - SPT-bit set, J - Join SPT
```

```
Timers: Uptime/Expires
```

```
DVMRP (191.9.1.11/32, 234.1.1.1), 00:00:36/00:00:00, flags:
  Incoming interface: vlan.0.3200
  Outgoing interface list:
```

```
DVMRP (191.9.1.12/32, 234.1.1.1), 00:00:36/00:00:00, flags:
  Incoming interface: vlan.0.3200
  Outgoing interface list:
```

```
DVMRP (193.9.3.30/32, 234.3.3.3), 3d, 01:13:10/00:00:00, flags:
  Incoming interface: vlan.0.930
  Outgoing interface list:
    vlan.0.3100, Forward/DVMRP, 2d, 19:32:38/00:00:00
```

```
DVMRP (193.9.3.31/32, 234.3.3.3), 3d, 01:13:04/00:00:00, flags:
  Incoming interface: vlan.0.930
  Outgoing interface list:
    vlan.0.3100, Forward/DVMRP, 2d, 19:32:38/00:00:00
```

```
DVMRP (193.9.3.32/32, 234.3.3.3), 3d, 01:13:11/00:00:00, flags:
  Incoming interface: vlan.0.930
  Outgoing interface list:
    vlan.0.3100, Forward/DVMRP, 2d, 19:32:38/00:00:00
```

## Prune Messages

If a device receives a datagram that has no IGMP group members present, and all the downstream networks are leaf networks, the device sends a prune packet upstream to the source tree.

When sending a prune upstream, the device:

1. Decides if the upstream neighbor is capable of receiving prunes.
  - If it is not, then the sending device proceeds no further.
  - If it is, then the sending device proceeds as follows.
2. Stops any pending grafts awaiting acknowledgments.
3. Determines the prune lifetime.

This value should be the minimum of the default prune lifetime (randomized to prevent synchronization) and the remaining prune lifetimes of the downstream neighbors.

4. Forms and transmits the packet to the upstream neighbor for the source.

To ensure the prune is accepted, the DVMRP-enabled device sets a negative cache prune entry for three seconds. If the traffic has not stopped after three seconds, the device sends another prune and doubles the cache entry. This method is called exponential back-off. The more prunes that are dropped, the longer the back-off becomes.

After the prune lifetime expires (two hours), the prune transmission process is repeated.

When receiving a prune, the upstream device:

1. Decides if the sending neighbor is known.
  - If the neighbor is unknown, it discards the received prune.
  - If the neighbor is known, the receiving device proceeds as follows.
2. Ensures the prune message contains at least the correct amount of data.
3. Copies the source address, group address, and prune time-out value, and, if it is available in the packet, the netmask value to determine the route to which the prune applies.
4. Determines if there is active source information for the source network, multicast group (S,G) pair.
  - If there is not, then the device ignores the prune.
  - If there is, then the device proceeds as follows.
5. Verifies that the prune was received from a dependent neighbor for the source network.
  - If it was not, then the device discards the prune.
  - If it was, then the device proceeds as follows.
6. Determines if a prune is currently active from the same dependent neighbor for this S,G pair.
  - If not active, creates a state for the new prune and sets a timer for the prune lifetime
  - If active, resets the timer to the new time-out value.
7. Determines if all dependent downstream devices on the interface from which the prune was received have now sent prunes.
  - If they have not, removes the interface from all forwarding cache entries for this group instantiated using the route to which the prune applies.
  - If they have, determines if there are group members active on the interface and if this device is the designated forwarder for the network.

### Graft Messages

Leaf devices send graft messages when the following occur:

- A new local member joins a group that has been pruned upstream and this device is the designated forwarder for the source.
- A new dependent downstream device appears on a pruned branch.
- A dependent downstream device on a pruned branch restarts.
- A graft retransmission timer expires before a graft ACK is received.

Graft messages are sent upstream hop-by-hop until the multicast tree is reached. Since there is no way to tell whether a graft message was lost or the source has stopped sending, each graft message is acknowledged hop-by-hop.

When sending grafts, the downstream device does the following:

1. Verifies a prune exists for the source network and group.
2. Verifies that the upstream device is capable of receiving prunes (and therefore grafts).




3. Adds the graft to the retransmission timer list awaiting an acknowledgment.
4. Formulates and transmits the graft packet.

When receiving grafts, the upstream device does the following:

1. Verifies whether the neighbor is known.
  - If unknown, discards the received graft.
  - If known, proceeds as follows.
2. Ensures the graft message contains at least the correct amount of data.
3. Sends back a graft ACK to the sender.
4. If the sender was a downstream dependent neighbor from which a prune had previously been received:
  - Removes the prune state for this neighbor.
  - If necessary, updates any forwarding cache entries based on this (source, group) pair to include this downstream interface.

Figure 23-3 shows the DVMRP pruning and grafting process.

**Figure 23-3 DVMRP Pruning and Grafting**



# Protocol Independent Multicast (PIM)

## Overview



**Note:** PIM requires an advanced router license. See “[License Overview](#)” on page 9-10 for 7100-Series advanced router license details.

PIM is made up of a collection of multicast routing protocols optimized for different multicast environments. The 7100-Series platform supports two PIM modes: sparse mode (SM) and dense mode (DM).

PIM dynamically builds a distribution tree for forwarding multicast data on a network. It is designed for use where there may be many devices communicating at the same time, and any one of the devices could be the sender at any particular time. Scenarios for using PIM multicasting include desktop video conferencing and telephone conference calls.

PIM relies on IGMP technology to determine group memberships and uses existing unicast routes to perform reverse path forwarding (RPF) checks, which are, essentially, a route lookup on the source. Its routing engine then returns the best interface, regardless of how the routing table is constructed. In this sense, PIM is independent of any routing protocol. It can perform RPF checks using protocol-specific routes (for example, OSPF routes), static routes, or a combination of route types.



**Note:** IGMP must be enabled for PIM to operate.

PIM-SM is a multicast routing protocol optimized for a network in which multicast recipients are sparsely distributed throughout the network. PIM-SM assumes that the network contains subnets that will not want a given multicast packet. Given this assumption, routers must explicitly tell their upstream neighbors about their interest in particular groups and sources. PIM-SM creates multicast distribution trees based upon the sending and receiving of PIM Join and Prune messages to join or leave the distribution tree.

PIM-SM by default uses shared trees, which are multicast distribution trees rooted at some selected node called the Rendezvous Point (RP) and used by all sources sending to the multicast group. PIM-SM is a soft-state protocol. All state periodically times out after receiving the control message that instantiated it. To keep the state alive, all PIM Join messages are periodically retransmitted.


PIM-DM is a multicast routing protocol optimized for a network in which the receivers for any multicast group are assumed to be densely distributed throughout the network. PIM-DM assumes that most, or at the very least, many subnets in the network will want any given multicast packet. PIM-DM uses source-based trees rather than the shared (RP-based) tree used by PIM-SM. PIM-DM does not use the concept of an RP in which all sources first send packets to a single router. PIM-DM initially sends multicast data to all hosts in the network. Routers that do not have any interested hosts then send PIM Prune messages to remove themselves from the tree. Because PIM-DM does not use RPs, it is easier to implement and deploy than PIM-SM. It is an efficient protocol when most receivers are interested in the multicast data, but does not scale well across larger domains in which most receivers are not interested in the data.

## PIM-SM

PIM-SM, a shared distribution tree technology, designates a router as the rendezvous point (RP), which is the root of a shared tree for a particular group. All sources send packets to the group via the RP (that is, traffic flows from the sender to the RP, and from the RP to the receiver). By maintaining one RP-rooted tree instead of multiple source-rooted trees, bandwidth is conserved.

Figure 23-4 illustrates the PIM traffic flow.

**Figure 23-4 PIM Traffic Flow**



1. The source's DR registers (encapsulates) and sends multicast data from the source directly to the RP via a unicast routing protocol (number 1 in figure). The RP de-encapsulates each register message and sends the resulting multicast packet down the shared tree.
2. The last-hop router (that is, the receiver's DR) sends a multicast group (\*,G) join message upstream to the RP, indicating that the receiver wants to receive the multicast data (number 2 in figure). This builds the RP tree (RPT) between the last-hop router and the RP.
3. The RP sends an S,G join message to the source (number 3 in figure). It may send the join message immediately, or after the data rate exceeds a configured threshold. This allows the administrator to control how PIM-SM uses network resources.
4. The last-hop router joins the shortest path tree (SPT) and sends an S,G join message to the source. (number 4 in figure). This builds the SPT.
5. Native multicast packets (that is, non-registered packets) are sent from the source's DR to the receiver on its SPT (number 5 in figure), while registered multicast packets continue to be sent from the source's DR to the RP.
6. A prune message is sent from the last-hop router to the RP (number 6 in figure).
7. A prune message (*register-stop*) is sent from the RP to the source's DR (number 7 in figure). Once traffic is flowing down the SPT, the RPT is pruned for that given S,G.

When receivers go away, prunes are sent (S,G prune messages towards the source on the SPT, and \*,G prune messages towards the RP on the RPT). When new receivers appear, the process begins again.

## PIM Support on Extreme Networks Devices

Extreme Networks devices support version 2 of the PIM protocol as described in RFC 4601 and *draft-ietf-pim-sm-v2-new-09*.

The PIM specifications define several modes or methods by which a PIM router can build the distribution tree. Extreme Networks devices support sparse mode (PIM-SM) and source-specific multicast (PIM-SSM).

PIM-SM uses only those routers that need to be included in forwarding multicast data. PIM-SM uses a host-initiated process to build and maintain the multicast distribution tree. Sparse mode routers use bandwidth more efficiently than other modes, but can require more processing time when working with large numbers of streams.

PIM-SSM is a subset of the PIM-SM protocol. PIM-SSM is disabled by default and must be explicitly enabled. PIM-SSM builds trees that are rooted in just one source, offering a more secure and scalable model for a limited amount of applications such as broadcasting of content. PIM-SSM is not independent of PIM-SM. PIM-SM must be enabled on all interfaces that use PIM-SSM. In PIM-SSM, an IP datagram is transmitted by a source S to an SSM destination address G, and receivers can receive this datagram by subscribing to channel (S,G). The destination address range for PIM SSM is 232.0.0.0/8 for IPv4 and ff3x:0000/32 where (x = 4,5,8, or E) for IPv6.

PIM-SSM does not require an RP candidate or BSR candidate. In a mixed PIM-SM and PIM-SSM configuration, the RP candidate and BSR candidate need to be configured for the PIM-SM group address range only. Enable IGMP on all PIM-SSM interfaces and enable IGMP querying on the PIM-SSM receiver interface. PIM-SSM requires IGMPv3 and MLDv2 at the edge of the network to process the source-specific IGMP and MLD joins.

### Key Features

Key features of PIM-SM are the following:

- Uses IGMP to propagate group membership information
- Sends hello messages to determine neighbor presence and configuration
- Sends join/prune messages to determine the need to retain multicast route information for a particular group on an interface
- Sends assert messages to resolve conflicts that occur regarding inbound interfaces
- Uses routes in the Multicast Routing Information Base (MRIB) to perform its reverse path forwarding check

Key features of PIM-SSM are the following:

- Protects against Denial of Service Attacks from unwanted sources
- Is easier to provision and maintain due to the single source address that a receiver can request data from
- Provides the ideal mechanism for internet broadcasts that originate from a single source and go to multiple receivers
- Does not require unique multicast addresses; it depends upon the receiver request for the destination address of the broadcast

Key features of PIM-DM are the following:

- Ease of configuration
- Operational and overhead efficiencies when a high density of network PIM enabled router interfaces have receivers attached for a given multicast data stream

## PIM-SM Message Types

Extreme Networks PIM-SM-enabled devices use the following message types:

- Hello — These messages announce the sender's presence to other PIM-SM devices. The hello packet includes options such as:
  - Hold time — the length of time to keep the sender reachable
  - Designated router (DR) priority — used to designate which PIM-SM device will act on behalf of sources and receivers in the PIM-SM domain
- Register — These messages are used by a source's DR to encapsulate (register) multicast data, and send it to the rendezvous point (RP) — a PIM-SM router designated as the root of a shared tree.
- Register-Stop — These messages are used by the RP to tell the source's DR to stop registering traffic for a particular source.
- Join/Prune (J/P) — These messages contain information on group membership received from downstream routers.

PIM-SM adopts RPF technology in the join/prune process. When a multicast packet arrives, the router first judges the correctness of the arriving interfaces:

- If the packet is a source address/multicast group (S,G) entry (on the shortest path tree (SPT)), then the correct interface is the reverse path forwarding (RPF) interface towards the source.
- If the packet is not an S,G entry (on the RP tree (RPT)), then the correct interface is the RPF interface towards the RP.

A router directly connected to the hosts is often referred to as a leaf router or DR. The leaf router is responsible for sending the prune messages to the RP, informing it to stop sending multicast packets associated with a specific multicast group. When the RP receives the prune message, it will no longer forward the multicast traffic out the interface on which it received the prune message.

- Assert — These messages indicate that the device received a data packet on its outbound (receiving) interface for the group. They report the metric or distance to the source or RP to help the device identify the most direct path to the root of the tree. If multiple routers claim to have the most direct path to the source or RP, each device sends its own assert message and the router with the best metric wins. The other device will then remove that link from its outbound interface list for the group.
- Bootstrap — These messages are sent by the PIM-SM router that has been elected as the bootstrap router (BSR) to inform all PIM-SM routes of the RP/group mappings.
- Candidate RP message — These messages are sent by the configured candidate RP routers to the BSR to inform the BSR of its RP/group candidacy.

## PIM-SSM Message Types

The PIM-SSM implementation is a subset of PIM-SM protocol. PIM-SM and PIM-SSM can coexist on a single router and are both implemented using the PIM-SM protocol.

Extreme Networks PIM-SSM enabled devices use the following PIM-SM message types:

- Hello — These messages announce the sender's presence to other PIM-SM devices. The hello packet includes options such as:
  - Hold time — the length of time to keep the sender reachable
  - Designated router (DR) priority — used to designate which PIM-SM device will act on behalf of sources and receivers in the PIM-SM domain

- Join/Prune (J/P) — These messages contain information on group membership received from downstream routers.
- PIM-SM adopts RPF technology in the join/prune process. When a multicast packet arrives, the router first judges the correctness of the arriving interfaces:
  - If the packet is a source address/multicast group (S,G) entry (on the shortest path tree (SPT)), then the correct interface is the reverse path forwarding (RPF) interface towards the source.
- Assert — These messages indicate that the device received a data packet on its outbound (receiving) interface for the group. They report the metric or distance to the source to help the device identify the most direct path to the root of the tree. If multiple routers claim to have the most direct path to the source, each device sends its own assert message and the router with the best metric wins. The other device will then remove that link from its outbound interface list for the group.

### PIM-DM Message Types

The PIM-DM-enabled devices use the following message types:

- Hello — These messages announce the sender's presence to other PIM-DM devices. The hello packet includes options such as:
  - Hold time — the length of time to keep the sender reachable
- Join/Prune (J/P) — These messages contain information on group membership received from downstream routers.
- Graft — These messages are sent upstream when new group membership is added to a pruned branch, instructing the upstream router to forward multicast data for the specified source to the downstream router.
- Graft ACK — These messages are sent to the downstream router, acknowledging the reception of a graft message from the downstream router.
- State Refresh — These messages are generated periodically by the PIM-DM router directly connected to a source and sent to neighbor routers. State refresh message minimize network overhead by conveying prune state.

### Anycast-RP

The 7100-Series supports anycast-RP. Anycast-RP provides a means of avoiding a single point of failure through fast convergence when a PIM RP router fails. It also provides for RP load balancing. The relationship between a source or receiver and the PIM RP router is a one-to-one relationship. The relationship between a source or receiver and an Anycast-RP set of routers is a one-to-many relationship, where one of multiple anycast configured RPs is selected by the routing protocol to be the source or receiver PIM RP router.

Anycast-RP provides for the selection of a set of routers to be identified as anycast RPs by

- Configuring each member of the anycast-RP set as either a static RP or a PIM candidate RP using the same loopback anycast IP address as the RP IP address
- Configuring:
  - A loopback interface with the same IP address for each anycast-RP router in the set
  - Either a second loopback interface or another hardware interface to be configured with a unique address for this peer of the anycast-RP set

Each anycast-RP router is configured with the same anycast-RP address and all the peer-addresses of each router in the anycast-RP router set. A unique peer address is used to allow each member of the anycast-RP set to identify all other members of the set. Each anycast-RP and peer-address combination is configured in its own command line entry using the **ip pim anycast-rp** command.

The routing protocol determines which member of the anycast-RP router set will function as the PIM RP router. Should the PIM RP router fail, the routing protocol determines the next anycast-RP router that will become the new PIM RP router, based upon the routing protocol's routing criteria.

[Figure 23-5](#) on page 23-17 illustrates an Anycast-RP configuration example.

### RP1

- Create and enable VLAN 10 with IP interfaces
- Configure the underlying unicast routing protocol (OSPF)
- Enable IGMP on VLAN 10
- Configure interface loopback 1 with the anycast-RP address 1.1.1.1/32
- Configure interface loopback 2 with the peer-address 10.0.0.1/32
- Configure 1.1.1.1 as either a static RP using the **ip pim rp-address** command or an RP candidate using the **ip pim rp-candidate** command
- Configure RP 1.1.1.1 as an anycast-RP set with the peer-addresses for RP1, RP2, and RP3 using the following commands:
  - **ip pim anycast-rp 1.1.1.1 10.0.0.1**
  - **ip pim anycast-rp 1.1.1.1 20.0.0.1**
  - **ip pim anycast-rp 1.1.1.1 30.0.0.1**

### RP2


- Create and enable VLAN 20 with IP interfaces
- Configure the underlying unicast routing protocol (OSPF)
- Enable IGMP on VLAN 20
- Configure interface loopback 1 with the anycast-RP address 1.1.1.1/32
- Configure interface loopback 2 with the peer-address 20.0.0.1/32
- Configure 1.1.1.1 as either a static RP using the **ip pim rp-address** command or an RP candidate using the **ip pim rp-candidate** command
- Configure RP 1.1.1.1 as an anycast-RP set with the peer-addresses for RP1, RP2, and RP3 using the following commands:
  - **ip pim anycast-rp 1.1.1.1 10.0.0.1**
  - **ip pim anycast-rp 1.1.1.1 20.0.0.1**
  - **ip pim anycast-rp 1.1.1.1 30.0.0.1**

### RP3

- Create and enable VLAN 30 with IP interfaces
- Configure the underlying unicast routing protocol (OSPF)
- Enable IGMP on VLAN 30
- Configure interface loopback 1 with the anycast-RP address 1.1.1.1/32

- Configure interface loopback 2 with the peer-address 30.0.0.1/32
- Configure 1.1.1.1 as either a static RP using the `ip pim rp-address` command or an RP candidate using the `ip pim rp-candidate` command
- Configure RP 1.1.1.1 as an anycast-RP set with the peer-addresses for RP1, RP2, and RP3 using the following commands:
  - `ip pim anycast-rp 1.1.1.1 10.0.0.1`
  - `ip pim anycast-rp 1.1.1.1 20.0.0.1`
  - `ip pim anycast-rp 1.1.1.1 30.0.0.1`

**Figure 23-5 Anycast-RP Configuration**



With all anycast-RPs configured, the routing protocol selects RP3 as the RP for this domain based upon its routing criteria. Should RP3 fail, the routing protocol will determine which of the remaining routers in the anycast-RP set will take over as RP. Should the failed router return to an operational state, the routing protocol will determine whether a new PIM RP will be selected based upon current conditions.

## PIM-DM

PIM-DM is a source-based tree technology. Multicast data is sent to all hosts in the network. Routers that do not have any interested hosts send PIM Prune messages to remove themselves from the tree associated with the multicast data source. If a router connects to a LAN with



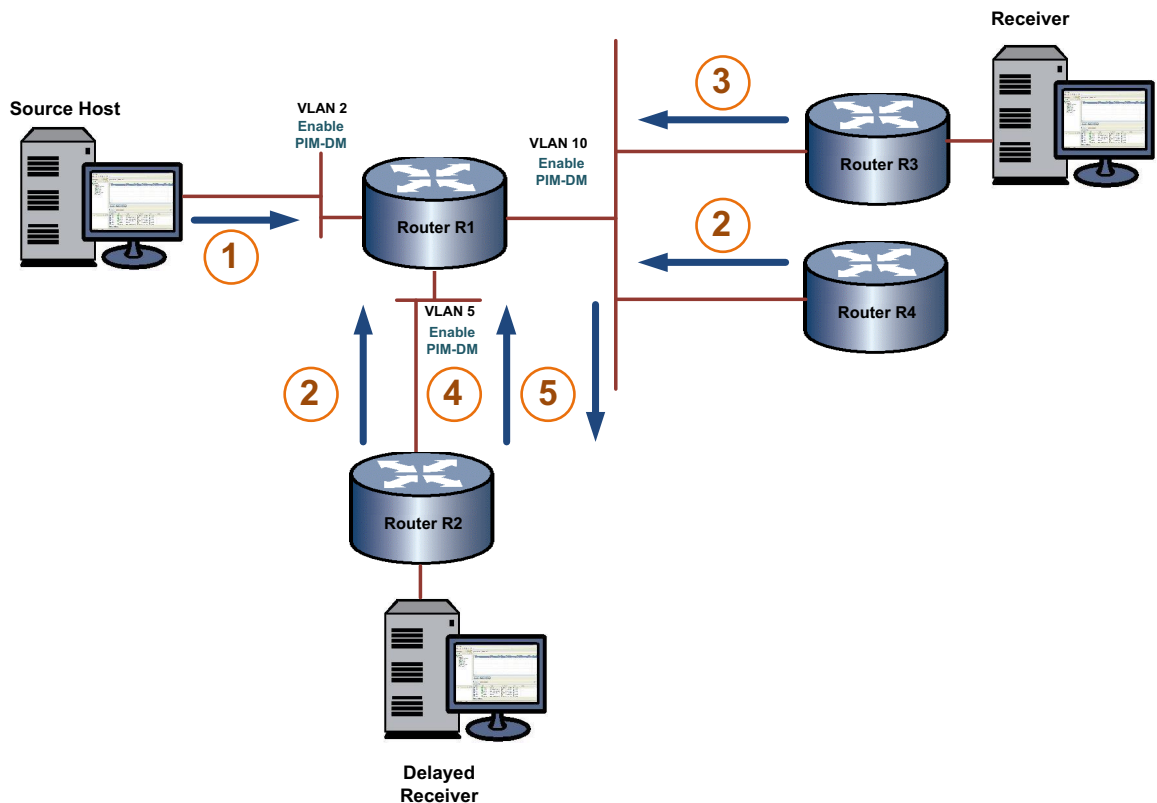
multiple routers, and that LAN contains both non-interested and interested routers for a given multicast data source:

- Non-interested routers will send a prune message upstream towards the source to inform the sending router that it does not have any receivers for the data stream
- Interested routers, having received a prune message, will send a join message upstream towards the source to inform the sending router that it has receivers for the data stream

Should a non-interested router become interested in a given multicast data stream, a graft message is sent upstream towards the source, informing the upstream router of its interest. The upstream router returns a graft ACK message acknowledging reception of the graft message. If no acknowledgment is received prior to the expiration of the graft retry timer (GRT), the interested router sends another graft message upstream.

Figure 23-6 illustrates the PIM-DM traffic flow.

**Figure 23-6 PIM-DM Traffic Flow**



1. The source sends multicast data on VLAN 2. R1 floods the data out all interfaces enabled for PIM-DM.
2. When R2 and R4 receive the multicast data, at this time no downstream receivers exist for the data's multicast group. At this time, both R2 and R4 send prune messages on the upstream interfaces indicating they have no receivers for the source data.
3. R3 receives the prune sent by R4 and responds by sending a join out its upstream interface indicating that it does have a receiver for the source data.
4. A receiver connected to R2 joins the multicast group for the source data. R2 responds by sending a graft message out its upstream interface informing R1 that it should now send its source data out its downstream interfaces to R2.

5. When R1 receives the graft message, a graft ACK message is returned to R2. R2 initiates a graft retry timer upon sending the graft message to R1. If R2 does not receive a graft ACK message from R1 before the timer expires, a new graft message is sent to R1.

When receivers go away, prunes are sent (S,G prune messages are sent towards the source). When new receivers appear, the process begins again.

## PIM Terms and Definitions

Table 23-1 lists terms and definitions used in PIM configuration.

**Table 23-1 PIM Terms and Definitions**

Term	Definition
Bootstrap Router (BSR)	<p>A PIM router responsible for collecting, within a PIM domain, the set of potential rendezvous points (RPs) and distributing the RP set information to all PIM routers within the domain. The BSR is dynamically elected from the set of candidate BSRs.</p> <p>RP set information includes group-to-RP mappings.</p>
Candidate Bootstrap Router (Candidate-BSR)	<p>A small number of routers within a PIM domain are configured as candidate BSRs, and each C-BSR is given a BSR priority. All C-BSRs multicast bootstrap messages (BSMs) containing their priority to the ALL-PIM-ROUTERS group. When a C-BSR receives a bootstrap message from a C-BSR with a higher priority, it stops sending. This continues until only one C-BSR remains sending bootstrap messages, and it becomes the elected BSR for the domain.</p>
Rendezvous Point (RP)	<p>The root of a group-specific distribution tree whose branches extend to all nodes in the PIM domain that want to receive traffic sent to the group.</p> <p>RPs provide a place for receivers and senders to meet. Senders use RPs to announce their existence, and receivers use RPs to learn about new senders of a group.</p> <p>The RP router, for the group, is selected by using the hash algorithm defined in RFC 2362.</p>
Candidate Rendezvous Point (Candidate-RP)	<p>PIM routers configured to participate as RPs for some or all groups.</p> <p>C-RPs send C-RP Advertisement messages to the BSR. The messages contain the list of group prefixes for which the C-RP is willing to be the RP. Once the PIM-SM routers receive the BSR's message, the routers use a common hashing algorithm to hash the C-RP address, group, and mask together to identify which router will be the RP for a given group.</p> <p>A C-RP router must also learn which PIM-SM router is the BSR. Each designated candidate-BSR (C-BSR) asserts itself as the BSR, then defers once it receives a preferable BSR message. Eventually, all C-RPs send their messages to a single BSR, which communicates the <i>Candidate RP-set</i> to all PIM-SM routers in the domain.</p>
dense mode	<p>PIM dense mode (DM) uses a source-based tree to distribute multicast data. DM does not require routers on the network to explicitly request interest in a given data stream and assumes that most routers in any given network will be interested in order to maintain efficient operation.</p>
Static RP	<p>If a BSR is not used to distribute RP set information, RP-to-group mappings are configured statically on each router.</p> <p>Static RP configuration and use of bootstrap routers are mutually exclusive. You should not configure both in a PIM-SM domain because such configuration could result in inconsistent RP sets. Statically configured RP set information will take precedence over RP set information learned from a BSR.</p>

**Table 23-1 PIM Terms and Definitions (continued)**

Term	Definition
Anycast-RP	Anycast-RP provides a means of fast convergence when a PIM RP router fails. All members of the anycast-RP set share the same IP address configured on a loopback interface of each set member. A peer-address associated with the member specifies a unique IP address that identifies the router and can be either a loopback or physical interface.
Designated Router (DR)	A designated router is elected from all the PIM routers on a shared network. DRs are responsible for encapsulating multicast data from local sources into PIM-SM register messages and for unicasting them to the RP. The router with the highest priority wins the DR election. In the case of a tie, the router with the highest IP address wins.
PIM Domain	A contiguous set of routers that implement PIM and are configured to operate within a common boundary defined by PIM multicast border routers.
PIM Multicast Border Router (PMBR)	A router that connects a PIM domain to other multicast routing domains.
sparse mode	PIM sparse mode (SM) uses a host-initiated process to build and maintain the multicast distribution tree, using only those routers that need to be included in forwarding multicast data. Sparse mode routers use bandwidth more efficiently than other modes, but can require more processing time when working with large numbers of streams
source-specific multicast	PIM source-specific multicast (SSM) is a modular switch only subset of the PIM-SM protocol that builds trees rooted in just one source and is used by applications such as content broadcasting.

## Configuring Multicast

This section provides the following information about configuring multicast.

For information about...	Refer to page...
<a href="#">Configuring IGMP</a>	<a href="#">23-20</a>
<a href="#">Configuring DVMRP</a>	<a href="#">23-23</a>
<a href="#">Configuring PIM</a>	<a href="#">23-25</a>

## Configuring IGMP

IGMP is configured at the switch level in any command mode on the 7100-Series devices. At Layer 2, IGMP can be enabled for VLANs, regardless of whether it is enabled on routed interfaces. If, however, IGMP is enabled on a routed interface, and the routed interface is a routed VLAN, then IGMP must also be enabled at the switch level.

### IGMP Configuration Commands

[Table 23-2](#) lists the IGMP configuration commands for 7100-Series devices.

**Table 23-2 IGMP Configuration Commands**

Task	Command
Enable IGMP on one or more VLANs.	<b>set igmp enable <i>vlan-list</i></b>

**Table 23-2 IGMP Configuration Commands (continued)**

Task	Command
Disable IGMP on one or more VLANs.	<b>set igmp disable</b> <i>vlan-list</i>
Enable IGMP querying on one or more VLANs.	<b>set igmp query-enable</b> <i>vlan-list</i>
Disable IGMP querying on one or more VLANs.	<b>set igmp query-disable</b> <i>vlan-list</i>
Determine what action to take with multicast frames when the multicast group table is full.	<b>set igmp grp-full-action</b> <i>action</i>
Configure IGMP settings on one or more VLANs.	<b>set igmp config</b> <i>vlan-list</i> {[ <b>query-interval</b> <i>query-interval</i> ] [ <b>igmp-version</b> <i>igmpversion</i> ] [ <b>max-resp-time</b> <i>max-resp-time</i> ] [ <b>robustness</b> <i>robustness</i> ] [ <b>last-mem-int</b> <i>last-mem-int</i> ] [ <b>fast-leave</b> <i>fast-leave</i> ] [ <b>rtr-alert-checking</b> <i>rtr-alert-checking</i> ] [ <b>filter-id</b> <i>filter-id</i> ] [ <b>filter-status</b> {enable   disable}]}
Remove IGMP configuration settings for one or more VLANs.	<b>set igmp delete</b> <i>vlan-list</i>
Change the IGMP classification of received IP frames.	<b>set igmp protocols</b> [ <b>classification</b> <i>classification</i> ] [ <b>protocol-id</b> <i>protocol-id</i> ] [ <b>modify</b> ]
Clear the binding of IP protocol ID to IGMP classification.	<b>clear igmp protocols</b> [ <b>protocol-id</b> <i>protocol-id</i> ]
Creates a new static IGMP entry or to adds one or more new include or exclude ports to an existing entry.	<b>set igmp static</b> <i>group</i> <i>vlan-list</i> [ <b>modify</b> ] [ <b>include-ports</b> <i>include-ports</i> ] [ <b>exclude-ports</b> <i>exclude-ports</i> ]
Create an input filter to apply to the VLAN.	<b>set igmp input-filter</b> <i>filter-id</i> <i>rule-id</i> <b>start-ip</b> <i>ip-address</i> <b>end-ip</b> <i>ip-address</i> <b>protocol-action</b> {deny   allow} <b>flow-action</b> {drop   flood   allow}
Clear an input filter.	<b>clear igmp input-filter</b> <i>filter-id</i> [ <i>rule-id</i> ]
Set the action taken when the first few frames of a multicast stream are received (that is, before the stream is added to the MLD database).	<b>set igmp unknown-input-action</b> {routers   flood   discard}

## Basic IGMP Configurations

[Procedure 23-1](#) describes the basic steps to configure IGMP on 7100-Series devices. This procedure assumes that the VLANs on which IGMP will run have been configured and enabled with IP interfaces.

### Procedure 23-1 Basic IGMP Configuration

Step	Task	Command
1.	In switch mode, configure IGMP for each VLAN interface.	<b>set igmp config</b> <i>vlan-list</i> {[ <b>query-interval</b> <i>query-interval</i> ] [ <b>igmp-version</b> <i>igmpversion</i> ] [ <b>max-resp-time</b> <i>max-resp-time</i> ] [ <b>robustness</b> <i>robustness</i> ] [ <b>last-mem-int</b> <i>last-mem-int</i> ] [ <b>fast-leave</b> <i>fast-leave</i> ] [ <b>rtr-alert-checking</b> <i>rtr-alert-checking</i> ] [ <b>filter-id</b> <i>filter-id</i> ] [ <b>filter-status</b> {enable   disable}]}
2.	In switch mode, enable IGMP on each VLAN interface.	<b>set igmp enable</b> <i>vlan-list</i>

**Procedure 23-1 Basic IGMP Configuration**

Step	Task	Command
3.	In switch mode, enable IGMP querying on each of the VLANs specified in step 2.	<b>set igmp query-enable</b> <i>vlan-list</i>

For more information on IGMP CLI commands, refer to your device's *CLI Reference Guide*.

**Example IGMP Configuration**

```
7100 Chassis(su)->set igmp enable 2, 3
7100 Chassis(su)->set igmp query-enable 2, 3
```

**IGMP Display Commands**

[Table 23-3](#) lists Layer 2 IGMP show commands for 7100-Series devices.

**Table 23-3 Layer 2 IGMP Show Commands**

Task	Command
Display the status of IGMP on one or more VLANs.	<b>show igmp enable</b> <i>vlan-list</i>
Display the IGMP query status of one or more VLANs.	<b>show igmp query</b> <i>vlan-list</i>
Display the action to be taken with multicast frames when the multicast IGMP flow table is full.	<b>show igmp flow-full-action</b>
Display IGMP configuration information for one or more VLANs.	<b>show igmp config</b> <i>vlan-list</i>
Display IGMP information regarding multicast group membership.	<b>show igmp groups</b> [ <b>group</b> <i>group</i> ] [ <b>vlan-list</b> <i>vlan-list</i> ] [ <b>sip</b> <i>sip</i> ] [ <b>-verbose</b> ]
Display static IGMP ports for one or more VLANs or IGMP groups.	<b>show igmp static</b> <i>vlan-list</i> [ <b>group</b> <i>group</i> ]
Display the binding of IP protocol id to IGMP classification.	<b>show igmp protocols</b>
Display IGMP information for a specific VLAN.	<b>show igmp vlan</b> [ <i>vlan-list</i> ]
Display IGMP reporter information.	<b>show igmp reporters</b> [ <b>portlist</b> <i>portlist</i> ] [ <b>group</b> <i>group</i> ] [ <b>vlan-list</b> <i>vlan-list</i> ] [ <b>sip</b> <i>sip</i> ]
Display IGMP flow information.	<b>show igmp flows</b> [ <b>portlist</b> <i>portlist</i> ] [ <b>group</b> <i>group</i> ] [ <b>vlan-list</b> <i>vlan-list</i> ] [ <b>sip</b> <i>sip</i> ]
Display IGMP counter information.	<b>show igmp counters</b>
Display configuration information for input filters.	<b>show igmp input-filter</b> [ <i>filter-id</i> ] [ <i>rule-id</i> ]
Display the action taken when the first frames of a multicast stream are received.	<b>show igmp unknown-input-action</b>

[Table 23-4](#) lists Layer 3 IGMP show commands for 7100-Series devices.

**Table 23-4 Layer 3 IGMP Show Commands**

Task	Command
Display IGMP information regarding multicast group membership.	<b>show ip igmp groups</b>

**Table 23-4 Layer 3 IGMP Show Commands**

Task	Command
Display multicast-related information about a specific interface or all interfaces.	<b>show ip igmp interface [vlan <i>vlan-id</i>]</b>

## Configuring DVMRP

### DVMRP Configuration Commands

[Table 23-5](#) lists the DVMRP configuration commands for 7100-Series devices.

**Table 23-5 DVMRP Configuration Commands**

Task	Command
Enable or disable DVMRP on an interface.	<b>ip dvmrp</b> <b>no ip dvmrp</b>
Configure the metric associated with a set of destinations for DVMRP reports.	<b>ip dvmrp metric <i>metric</i></b>

### Basic DVMRP Configuration

By default, DVMRP is disabled globally on Extreme Networks 7100-Series devices and attached interfaces.

[Procedure 23-2](#) describes the basic steps to configure DVMRP on 7100-Series devices. This procedure assumes that the VLANs on which DVMRP will run have been configured and enabled with IP interfaces.


**Procedure 23-2 Basic DVMRP Configuration**

Step	Task	Command
1.	Configure IGMP for each VLAN interface.	<b>set igmp config <i>vlan-list</i> {[query-interval <i>query-interval</i>] [igmp-version <i>igmpversion</i>] [max-resp-time <i>max-resp-time</i>] [robustness <i>robustness</i>] [last-mem-int <i>last-mem-int</i>]}</b>
2.	Enable IGMP on each VLAN interface.	<b>set igmp enable <i>vlan-list</i></b>
3.	Enable DVMRP on each of the VLANs specified in step 2.	<b>ip dvmrp</b>

### Example DVMRP Configuration

[Figure 23-7](#) illustrates the DVMRP configuration of two 7100-Series devices shown in the example below. This example assumes the following:

- VLANs have been configured and enabled with IP interfaces
- IGMP has been enabled on the VLANs

**Figure 23-7 DVMRP Configuration on Two Routers**

### Router R1 Configuration

For the VLAN 1 interface, which provides connection to Router R2, an IP address is assigned and DVMRP is enabled. For the VLAN 2 interface, which provides connection to the host network, an IP address is assigned and DVMRP is enabled.

```
R1(su)->config
R1(su-config)->interface vlan 1
R1(su-config-intf-vlan.0.1)->ip address 192.0.1.2 255.255.255.0
R1(su-config-intf-vlan.0.1)->ip dvmrp
R1(su-config-intf-vlan.0.1)->no shutdown
R1(su-config-intf-vlan.0.1)->exit
R1(su-config)->interface vlan 2
R1(su-config-intf-vlan.0.2)->ip address 192.40.0.1 255.255.255.0
R1(su-config-intf-vlan.0.2)->ip dvmrp
R1(su-config-intf-vlan.0.2)->no shutdown
R1(su-config-intf-vlan.0.2)->exit
```

### Router R2 Configuration

For the VLAN 1 interface, which provides connection to the Router R1, an IP address is assigned and DVMRP is enabled. For the VLAN 3 interface which provides connection to the host network, an IP address is assigned and DVMRP is enabled.

```
R2(su)->config
R2(su-config)->interface vlan 1
R2(su-config-intf-vlan.0.1)->ip address 192.0.1.1 255.255.255.0
R2(su-config-intf-vlan.0.1)->ip dvmrp
R2(su-config-intf-vlan.0.1)->no shutdown
R2(su-config-intf-vlan.0.1)->exit
R2(su-config)->interface vlan 3
R2(su-config-intf-vlan.0.3)->ip address 192.20.0.1 255.255.255.0
R2(su-config-intf-vlan.0.3)->ip dvmrp
R2(su-config-intf-vlan.0.3)->no shutdown
R2(su-config-intf-vlan.0.3)->exit
```

## Displaying DVMRP Information

Table 23-6 lists the DVMRP show commands for 7100-Series devices.

**Table 23-6 DVMRP Show Commands**

Task	Command
Display information about the routes in the DVMRP routing table.	<b>show ip dvmrp route</b>

**Table 23-6 DVMRP Show Commands**

Task	Command
Display the IP multicast routing table.	<b>show ip mroute</b> [ <i>unicast-source-address</i>   <i>multicast-group-address</i> ] [ <b>summary</b> ]

Refer to the *Extreme Networks 7100-Series CLI Reference* for an example of each command's output.

## Configuring PIM

### PIM Configuration Commands

Table 23-7 lists the PIM configuration commands for Extreme Networks 7100-Series devices.

**Table 23-7 IPv4 PIM Sparse Mode Commands**

Task	Command
Enable PIM-SM on a routing interface. Use the <b>no</b> command to disable PIM-SM.	<b>ip pim sparse-mode</b> <b>no ip pim sparse-mode</b>
Enable PIM-SSM in router configuration mode. Use the <b>no</b> command to disable PIM-SSM.	<b>ip pim ssm</b> { <b>default</b>   <i>group-address group-mask</i> } <b>no ip pim ssm</b> { <b>default</b>   <i>group-address group-mask</i> }
Enable the router to announce its candidacy as a Bootstrap Router (BSR). Use the <b>no</b> command to remove the router as a BSR candidate.	<b>ip pim bsr-candidate</b> <i>pim-interface-address</i> [ <b>priority</b> <i>priority</i> ] <b>no ip pim bsr-candidate</b>
Set the priority for which a router will be elected as the designated router (DR). Use the <b>no</b> command to disable the DR functionality.	<b>ip pim dr-priority</b> <i>priority</i> <b>no ip dr-priority</b>
Set a static rendezvous point (RP) for a multicast group, specifying a specific group or a group-list. Use the <b>no</b> command to remove the static RP configuration.	<b>ip pim rp-address</b> <i>rp-address</i> { <i>group-address group-mask</i>   <b>group-list</b> <i>group-list</i> } <b>no ip rp-address</b> <i>rp-address</i> { <i>group-address group-mask</i>   <b>group-list</b> <i>group-list</i> }
Enable the router to advertise itself as a PIM candidate rendezvous point (RP) to the BSR specifying either a specific group or a group-list. Use the <b>no</b> command to remove the router as an RP candidate.	<b>ip pim rp-candidate</b> <i>pim-interface-address</i> { <i>group-address group-mask</i>   <b>priority</b> <i>priority</i>   <b>group-list</b> <i>group-list</i> [ <b>priority</b> <i>priority</i> ]} <b>no ip pim rp-candidate</b> <i>pim-interface-address</i> { <i>group-address group-mask</i>   <b>group-list</b> <i>group-list</i> [ <b>priority</b> <i>priority</i> ]}
Enable control of whether static RP configurations will override dynamic RP information learned for IPv4 groups.	<b>ip pim static-rp-override</b> <b>no ip pim static-rp-override</b>
Filter PIM neighbors by specifying a standard ACL containing neighbors to allow.	<b>ip pim neighbor-filter</b> <i>neighbor-filter</i> <b>no ip pim neighbor-filter</b> <i>neighbor-filter</i>
Configure an anycast Rendezvous Points (RP) set member for a multicast group. Use the <b>no</b> command to remove the specified anycast member.	<b>ip pim anycast-rp</b> <i>anycast-address peer-address</i> <b>no ip anycast-rp</b> <i>anycast-address peer-address</i>
Set the multicast graceful-restart period, which is the period of time in which a restarting router and its neighbors can continue to forward multicast packets during the failover.	<b>ip pim graceful-restart</b> [ <b>period</b> <i>value</i> ] <b>no ip pim graceful-restart</b>



**Table 23-7 IPv4 PIM Sparse Mode Commands**

Task	Command
Set PIM multicast to either load share over ECMP paths or have a single deterministic next hop for ECMP paths.	<b>ip pim multipath {hash   highest-nexthop}</b> <b>no ip pim multipath {hash   highest-nexthop}</b>

[Table 23-8](#) lists the PIM IPv6 sparse mode configuration commands for Extreme Networks 7100-Series devices.

**Table 23-8 IPv6 PIM Sparse Mode Commands**

Task	Command
Enable PIM-SM on a routing interface. Use the <b>no</b> command to disable PIM-SM.	<b>ipv6 pim sparse-mode</b> <b>no ipv6 pim sparse-mode</b>
Optionally enable PIM-SSM in router configuration mode. Use the <b>no</b> command to disable PIM-SSM	<b>ipv6 pim ssm {default   group-address/length}</b> <b>no ipv6 pim ssm {default   group-address/length}</b>
Enable the router to announce its candidacy as a Bootstrap Router (BSR). Use the <b>no</b> command to remove the router as a BSR candidate.	<b>ipv6 pim bsr candidate bsr interface-address [priority priority]</b> <b>no ipv6 pim bsr candidate bsr interface-address</b>
Set the priority for which a router will be elected as the designated router (DR). Use the <b>no</b> command to disable the DR functionality.	<b>ipv6 pim dr-priority priority</b> <b>no ipv6 pim dr-priority priority</b>
Set a static rendezvous point (RP) for a multicast group, specifying a specific group or a group-list. Use the <b>no</b> command to remove the static RP configuration.	<b>ipv6 pim rp-address rp-address group-list group-list</b> <b>no ipv6 pim rp-address rp-address group-list group-list</b>
Enable the router to advertise itself as a PIM candidate rendezvous point (RP) to the BSR specifying a group-list. Use the <b>no</b> command to remove the router as an RP candidate.	<b>ipv6 pim bsr candidate rp pim-interface-address {[group-list group-list] [priority priority]}</b> <b>no ipv6 pim bsr candidate bsr pim-interface-address {[group-list group-list] [priority priority]}</b>
Enable control of whether static RP configurations will override dynamic RP information learned for IPv6 groups.	<b>ipv6 pim static-rp-override</b> <b>no ipv6 pim static-rp-override</b>
Filter PIM neighbors by specifying a standard ACL containing neighbors to allow.	<b>ipv6 pim neighbor-filter neighbor-filter</b> <b>no ipv6 pim neighbor-filter neighbor-filter</b>
Configure an anycast Rendezvous Points (RP) set member for a multicast group. Use the <b>no</b> command to remove the specified anycast member.	<b>ipv6 pim anycast-rp anycast-address peer-address</b> <b>no ipv6 pim anycast-rp anycast-address peer-address</b>
Set the multicast graceful-restart period, which is the period of time in which a restarting router and its neighbors can continue to forward multicast packets during the failover.	<b>ipv6 pim graceful-restart [period value]</b> <b>no ipv6 pim graceful-restart</b>
Set PIM multicast to either load share over ECMP paths or have a single deterministic next hop for ECMP paths.	<b>ipv6 pim multipath {hash   highest-nexthop}</b> <b>no ipv6 pim multipath {hash   highest-nexthop}</b>

[Table 23-9](#) lists the PIM dense mode configuration commands for Extreme Networks 7100-Series devices.

**Table 23-9 PIM Dense Mode Commands**

Task	Command
Enable PIM-DM on a routing interface. Use the <b>no</b> command to disable PIM-DM.	<b>ip   ipv6 pim dense-mode</b> <b>no ip   ipv6 pim dense-mode</b>
Optionally set the interval between PIM dense mode state refresh messages.	<b>ip   ipv6 pim state-refresh origination-interval interval</b> <b>no ip   ipv6 pim state-refresh origination-interval interval</b>
Set PIM multicast to either load share over ECMP paths or have a single deterministic next hop for ECMP paths.	<b>ip   ipv6 pim multipath {hash   highest-nexthop}</b> <b>no ip   ipv6 pim multipath {hash   highest-nexthop}</b>

## Basic PIM-SM Configurations

The following describes a basic PIM configuration. PIM-SSM is a simplified version of PIM-SM. PIM-SSM does not require either a BSR or an RP. In a PIM-SSM configuration there is no need for a candidate BSR, a candidate RP, or a static RP. In a mixed PIM-SSM and PIM-SM configuration, the candidate BSR, candidate RP, and the static RP need only be configured for the non-PIM-SSM address ranges.

By default, PIM-SM and PIM-SSM are disabled globally on 7100-Series devices and attached interfaces. Basic PIM configuration includes the following steps:

1. Creating and enabling VLANs with IP interfaces.
2. Configuring the underlying unicast routing protocol (for example, OSPF).
3. Enabling IGMP on the VLANs. Enable IGMP for interfaces with IGMP reporters. Enable IGMP version 3 on interfaces using PIM-SSM
4. Configuring PIM-SM and/or PIM-SSM on the VLANs.

[Procedure 23-3](#), which describes the basic steps the configure PIM-SM on an 7100-Series device, assumes the following:

- VLANs have been configured and enabled with IP interfaces.
- The unicast routing protocol has been configured.
- IGMP has been enabled on the devices and VLANs that will be connected with hosts. For information on enabling IGMP, see [“Configuring IGMP”](#) on page 23-20.



**Note:** PIM-SSM and PIM-SM can coexist in a network. A candidate BSR, candidate RP, and static RP addresses can be configured in a PIM-SSM configuration, but are not required. Along with IGMP, PIM-SSM must be enabled on the source host interface and be reachable by the PIM-SSM destination addresses.

**Procedure 23-3 Basic PIM Sparse Mode Configuration**

Step	Task	Command(s)
1.	<p>If desired, change the DR priority of one or more interfaces on the Extreme Networks 7100-Series router from the default value of 1 in interface configuration mode.</p> <p>The highest priority PIM router on a shared network is elected the DR for that network.</p>	<p><b>IPv4:</b></p> <p><b>ip pim dr-priority</b> <i>priority</i></p> <p><b>IPv6:</b></p> <p><b>ipv6 pim dr-priority</b> <i>priority</i></p>
2.	<p>If the dynamic BSR RP set distribution method is used on the network, configure at least one PIM router as a candidate BSR in interface configuration mode.</p> <p>Note that the Extreme Networks 7100-Series router does not act as a BSR without being explicitly configured to do so.</p>	<p><b>IPv4:</b></p> <p><b>ip pim bsr-candidate</b> <i>pim-interface</i> [<b>priority</b> <i>priority</i>]</p> <p><b>IPv6:</b></p> <p><b>ipv6 pim bsr candidate bsr</b> <i>interface-address</i> [<b>priority</b> <i>priority</i>]</p>
3.	<p>If the dynamic BSR RP set distribution method will be used on the network, configure at least one PIM router as a Candidate Rendezvous Point in global configuration mode.</p> <p>Note that the Extreme Networks 7100-Series router does not act as an RP without being explicitly configured to do so.</p>	<p><b>IPv4:</b></p> <p><b>ip pim rp-candidate</b> <i>pim-interface</i> <i>group-address</i> <i>group-mask</i> [<b>priority</b> <i>priority</i>]</p> <p><b>IPv6:</b></p> <p><b>ipv6 pim bsr candidate rp</b> <i>pim-interface-address</i> {[<b>group-list</b> <i>group-list</i>] [<b>priority</b> <i>priority</i>]}</p>
4.	<p>If static RP set distribution is desired, configure the static RP set information in global configuration mode. The RP set information must be the same on all PIM routers in the network.</p> <p><b>Note:</b> Static RP set distribution cannot be combined with BSR RP set distribution in the same PIM domain. Routers with statically configured RP set information discard RP set information learned from a BSR.</p>	<p><b>IPv4:</b></p> <p><b>ip pim rp-address</b> <i>rp-address</i> <i>group-address</i> <i>group-mask</i></p> <p><b>IPv6:</b></p> <p><b>ipv6 pim rp-address</b> <i>rp-address</i> <b>group-list</b> <i>group-list</i></p>
5.	<p>Configure PIM-SM and/or PIM/SSM on the 7100-Series router that will run PIM-SM.</p> <p>PIM-SM is configured on the interface. PIM-SSM is globally configured in global configuration mode.</p> <p>IPv6 PIM-SSM is enabled on the device by default with an address range of FF3E:0000/32.</p>	<p><b>IPv4:</b></p> <p><b>ip pim sparse-mode</b></p> <p><b>ip pim ssm</b></p> <p><b>IPv6:</b></p> <p><b>ipv6 pim sparse-mode</b></p>

**PIM IPv4 and IPv6 Display Commands**

Table 23-10 lists the PIM IPv4 and IPv6 display commands for Extreme Networks 7100-Series devices.

**Table 23-10 PIM IPv4 and IPv6 Display Commands**

Task	Command
Display summary tables of PIM interfaces, neighbors, BSR, and group-to-RP mappings.	<b>show {ip   ipv6} pim</b>

**Table 23-10 PIM IPv4 and IPv6 Display Commands**

Task	Command
Display RP anycast information for all or a specified RP. (PIM-SM only).	<b>show {ip   ipv6} pim anycast-rp [rp-address rp-address]</b>
Display Bootstrap Router (BSR) information. (PIM-SM only).	<b>show {ip   ipv6} pim bsr [detail]</b>
Display information about PIM interfaces that are currently up (not shutdown).	<b>show {ip   ipv6} pim interface [ifName] [brief] [detail] [statistics]</b>
Display the PIM multicast route (*,G and S,G) table.	<b>show {ip   ipv6} pim mrt [source source   group group] [interface] [detail] [brief] [summary]</b>
Display the PIM multicast route (*,G and S,G) table by type.	<b>show {ip   ipv6} mrt type {all   s-g   star-g} [source source   group group] [interface] [detail] [brief] [type {all   s-g   star-g}] [summary]</b>
Display information about discovered PIM neighbors.	<b>show {ip   ipv6} pim neighbor [ifName] [brief] [detail] [statistics]</b>
Display the active rendezvous points (RPs) that are cached with associated multicast routing entries. (PIM-SM only).	<b>show {ip   ipv6} pim rp [mapping]</b>
Display the rendezvous point (RP) selected for a specified group. (PIM-SM only).	<b>show {ip   ipv6} pim rp-hash group-address</b>
Display PIM statistics for this device.	<b>show {ip   ipv6} pim statistics</b>
Display the multicast routing table.	<b>show {ip   ipv6} mroute [source source   group group   interface interface] [brief] [summary]</b>
Display the multicast forwarding cache that was used to program the hardware flow.	<b>show {ip   ipv6} mcache [group group   source source] [interface] [verbose   brief   summary] [statistics] [-wide]</b>


Refer to the *Extreme Networks 7100-Series CLI Reference* for a description of the output of each command.

## Example PIM Configuration

Figure 23-8 illustrates the PIM of four 7100-Series routers. For a PIM-DM configuration:

- Configure interfaces and enable IGMP as shown in the PIM-SM example script below
- Enable dense mode on all interfaces using the **ip pim dense-mode** command
- Optionally change the state refresh interval, using the **ip pim state-refresh origination-interval** command

The PIM-SM configuration is shown in the example scripts below. PIM-SM configuration includes configuring a preferred and a backup BSR for the topology, as well as two RPs for specific multicast groups and a backup RP for all groups.

**Figure 23-8 PIM-SM Configuration with Bootstrap Router and Candidate RPs**

### Router R1 Configuration

On this router, IGMP is enabled on VLAN 2, which connects to hosts, and PIM-SM is enabled on all interfaces. IGMP is used to determine host group membership on directly attached subnets. Note that IGMP is enabled in switch mode on 7100-Series routers.

VLAN 2 is configured as the backup candidate RP for all multicast groups by using the default RP priority of 192. Note that the C-RP with the smallest priority value is elected.

Alternatively, you could configure a loopback interface as a candidate RP, to avoid the dependency on a particular interface.

```
R1(su-config)->router id 1.1.1.1
R1(su-config)->interface vlan 2
R1(su-config-intf-vlan.0.2)->ip address 172.1.1.1 255.255.255.0
R1(su-config-intf-vlan.0.2)->no shutdown
R1(su-config-intf-vlan.0.2)->exit

R1(su)->set igmp enable 2
R1(su)->set igmp enable 3
R1(su)->set igmp enable 4
R1(su)->set igmp query-enable 2

R1(su-config)->ip pim rp-candidate 172.1.1.1 224.0.0.0 240.0.0.0
R1(su-config)->interface vlan 2
R1(su-config-intf-vlan.0.2)->ip pim sparse-mode
R1(su-config-intf-vlan.0.2)->exit
R1(su-config)->interface vlan 3
R1(su-config-intf-vlan.0.3)->ip address 172.1.2.1 255.255.255.0
R1(su-config-intf-vlan.0.3)->no shutdown
R1(su-config-intf-vlan.0.3)->ip pim sparse-mode
R1(su-config-intf-vlan.0.3)->exit
```

```
R1(su-config)->interface vlan 4
R1(su-config-intf-vlan.0.4)->ip address 172.1.3.1 255.255.255.0
R1(su-config-intf-vlan.0.4)->no shutdown
R1(su-config-intf-vlan.0.4)->ip pim sparse-mode
R1(su-config-intf-vlan.0.4)->exit
```

### Router R2 Configuration

On this router, PIM-SM is enabled on all interfaces. VLAN 9 is configured as a candidate BSR and is assigned a priority higher than the default of 0. Note that the C-BSR with the largest priority value is elected.

VLAN 9 is also configured as a candidate RP for the multicast group 224.2.2.0/24. Its priority is set to 2, which will most likely make it the elected RP for that particular group, since the C-RP with the smallest priority value is elected. (Note that Router R3 has an RP candidate priority value of 3 for that group.)

Again, alternatively, you could configure a loopback interface as a candidate BSR or RP, to avoid the dependency on a particular interface.

```
R2(su)->set igmp enable 3
R2(su)->set igmp enable 9
R1(su)->set igmp enable 8
R1(su)->set igmp enable 5
```

```
R2(su-config)->router id 1.1.1.2
R2(su-config)->ip pim bsr-candidate vlan 9 priority 2
```

```
R2(su-config)->interface vlan 3
R2(su-config-intf-vlan.0.3)->ip address 172.1.2.2 255.255.255.0
R2(su-config-intf-vlan.0.3)->no shutdown
R2(su-config-intf-vlan.0.3)->ip pim sparse-mode
R2(su-config-intf-vlan.0.3)->exit
```

```
R2(su-config)->interface vlan 9
R2(su-config-intf-vlan.0.9)->ip address 172.2.2.2 255.255.255.0
R2(su-config-intf-vlan.0.9)->no shutdown
R2(su-config-intf-vlan.0.9)->ip pim sparse-mode
R2(su-config-intf-vlan.0.9)->exit
R2(su-config)->ip pim rp-candidate 172.2.2.2 224.2.2.0 255.255.255.0 priority 2
```

```
R2(su-config)->interface vlan 8
R2(su-config-intf-vlan.0.8)->ip address 172.2.3.2 255.255.255.0
R2(su-config-intf-vlan.0.8)->no shutdown
R2(su-config-intf-vlan.0.8)->ip pim sparse-mode
R2(su-config-intf-vlan.0.8)->exit
```

```
R2(su-config)->interface vlan 5
R2(su-config-intf-vlan.0.5)->ip address 172.2.4.2 255.255.255.0
R2(su-config-intf-vlan.0.5)->no shutdown
R2(su-config-intf-vlan.0.5)->ip pim sparse-mode
R2(su-config-intf-vlan.0.5)->exit
```

### Router R3 Configuration

On this router, PIM-SM is enabled on all interfaces. VLAN 10 is configured as a backup candidate BSR, by leaving its priority at the default of 0.

VLAN 10 is also configured as a backup candidate RP for multicast group 224.2.2.0/24, by setting its priority value slightly higher (3) than the priority configured on R2 for the same group (2) (since the C-RP with the smallest priority value is elected).

```
R3(su)->set igmp enable 4
```

```
R3(su)->set igmp enable 8
R3(su)->set igmp enable 10
R3(su)->set igmp enable 6
R3(su)->configure

R3(su-config)->router id 1.1.1.3
R3(su-config)->ip pim bsr-candidate vlan 10

R3(su-config)->interface vlan 4
R3(su-config-intf-vlan.0.4)->ip address 172.1.3.3 255.255.255.0
R3(su-config-intf-vlan.0.4)->no shutdown
R3(su-config-intf-vlan.0.4)->ip pim sparse-mode
R3(su-config-intf-vlan.0.4)->exit

R3(su-config)->interface vlan 8
R3(su-config-intf-vlan.0.8)->ip address 172.2.3.3 255.255.255.0
R3(su-config-intf-vlan.0.8)->no shutdown
R3(su-config-intf-vlan.0.8)->ip pim sparse-mode
R3(su-config-intf-vlan.0.8)->exit

R3(su-config)->interface vlan 10
R3(su-config-intf-vlan.0.10)->ip address 172.3.3.3 255.255.255.0
R3(su-config-intf-vlan.0.10)->no shutdown
R3(su-config-intf-vlan.0.10)->ip pim sparse-mode
R3(su-config-intf-vlan.0.10)->exit
R3(su-config)->ip pim rp-candidate 172.3.3.3 224.2.2.0 255.255.255.0 priority 3

R3(su-config)->interface vlan 6
R3(su-config-intf-vlan.0.6)->ip address 172.3.4.3 255.255.255.0
R3(su-config-intf-vlan.0.6)->no shutdown
R3(su-config-intf-vlan.0.6)->ip pim sparse-mode
R3(su-config-intf-vlan.0.6)->exit
```

### Router R4 Configuration

This router does not play any special role in PIM-SM, except that it has hosts directly connected to it. IGMP is enabled on the interface that connects to hosts and PIM-SM is enabled on all interfaces.

```
R3(su)->set igmp enable 5
R3(su)->set igmp enable 6
R3(su)->set igmp enable 7
R3(su)->configure
R4(su-config)->router id 1.1.1.4
R4(su-config)#interface vlan 5
R4(su-config-intf-vlan.0.5)->ip address 172.2.4.4 255.255.255.0
R4(su-config-intf-vlan.0.5)->no shutdown
R4(su-config-intf-vlan.0.5)->ip pim sparse-mode
R4(su-config-intf-vlan.0.5)->exit

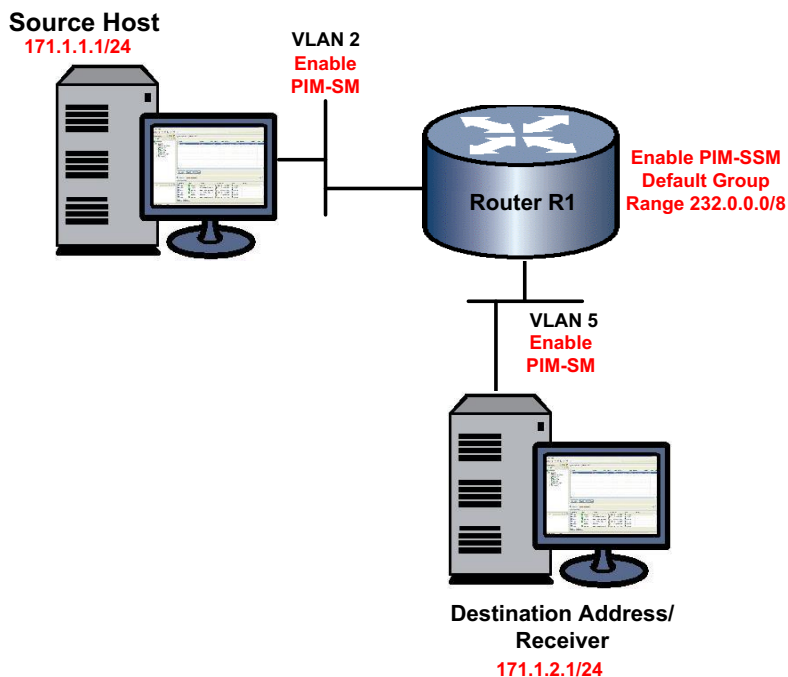
R4(su-config)->interface vlan 6
R4(su-config-intf-vlan.0.6)->ip address 172.3.4.4 255.255.255.0
R4(su-config-intf-vlan.0.6)->no shutdown
R4(su-config-intf-vlan.0.6)->ip pim sparse-mode
R4(su-config-intf-vlan.0.6)->exit

R4(su-config)->interface vlan 7
R4(su-config-intf-vlan.0.7)->ip address 172.4.4.4 255.255.255.0
R4(su-config-intf-vlan.0.7)->no shutdown
R4(su-config-intf-vlan.0.7)->ip pim sparse-mode
R4(su-config-intf-vlan.0.7)->exit
```

## Example PIM-SSM Configuration

Figure 23-9 illustrates the PIM-SSM configuration of a single router shown in the example scripts below. PIM-SSM is enabled on router R1 with the default group range of 232.0.0.0/8. VLANs connected to the source host and receiver are configured on the router. PIM-SSM and IGMP are enabled on all interfaces. IGMP query is enabled on the receiver interface.

Figure 23-9 PIM-SSM Configuration



### Router R1 Configuration

On this router:

- Enable PIM-SSM with the default group range
- Configure VLAN 2 with the source host IP address 171.1.1.1/24, and enable PIM-SSM on the interface
- Configure VLAN 5 with the receiver IP address 171.1.2.1/24, and enable PIM-SSM on the interface
- Enable IGMP version 3 on VLAN 2 and VLAN 5. IGMP is used to determine host group membership on directly attached subnets. PIM-SSM requires IGMP version 3. Note that IGMP version 2 is enabled by default in switch mode on 7100-Series routers.
- Enable IGMP querying on the receiver interface (VLAN 5)

```
R1(su-config)->router id 1.1.1.1
R1(su-config)->ip pim ssm default
R1(su-config)->interface vlan 2
R1(su-config-intf-vlan.0.2)->ip address 171.1.1.1 255.255.255.0
R1(su-config-intf-vlan.0.2)->ip pim sparse-mode
R1(su-config-intf-vlan.0.2)->no shutdown
R1(su-config-intf-vlan.0.2)->exit
R1(su-config)->interface vlan 5
R1(su-config-intf-vlan.0.5)->ip address 171.1.2.1 255.255.255.0
```



```
R1(su-config-intf-vlan.0.2)->ip pim sparse-mode
R1(su-config-intf-vlan.0.5)->no shutdown
R1(su-config-intf-vlan.0.5)->exit
R1(su-config)->exit
R1(su)->set igmp enable 2,5
R1(su)->set igmp query-enable 5
R1(su)->set igmp config 2,5 igmp-version 3
```

## Multicast Listener Discovery (MLD) Configuration

For information about...	Refer to page...
<a href="#">Using MLD in Your Network</a>	24-1
<a href="#">Implementing MLD</a>	24-1
<a href="#">Understanding MLD</a>	24-2
<a href="#">Configuring MLD</a>	24-5

### Using MLD in Your Network

Multicast Listener Discovery (MLD) enables each IPv6 router to discover the presence of nodes wishing to receive multicast packets on its directly attached links, and determines which multicast addresses are of interest to those neighboring nodes. MLD provides the discovered information to the routing protocol used by the router, in order to ensure that multicast packets are delivered to all links where there are interested receivers.

Multicast is a “one source to many destinations” method of simultaneously sending information over a network using the most efficient delivery strategy over each link. Only the end stations that explicitly indicate a need to receive a given multicast stream will receive it.

Applications that take advantage of multicast include video conferencing, streaming video, corporate communications, distance learning, and distribution of software, stock quotes, and news. See “[Multicast Configuration](#)” on page 23-1 for a detailed multicast discussion.

The 7100-Series supports MLD version 1 (RFC 2710) and version 2 (RFC 3810). MLD defaults to version 2.

Unlike unicast and broadcast, multicast uses network infrastructure efficiently because only one copy of the source traffic is sent throughout the network, going only to interested receivers, minimizing the burden placed on the sender, network, and receiver. The routers in the network take care of replicating the packet, where necessary, to reach multiple receivers. If a router decides that there are no interested users downstream from itself, it prunes the stream back to the next router. Thus, unwanted streams are not sent to the pruned routers, saving bandwidth and preventing unwanted packets from being sent.

### Implementing MLD

You can implement the MLD multicast protocol on Extreme Networks devices using simple CLI commands as described in this chapter and the Multicast Listener Discovery (MLD) Commands chapter of the *Extreme Networks 7100-Series CLI Reference*. A basic configuration process involves the following tasks:

1. Configuring the VLANs and IP interfaces on which you want to transmit multicast.

2. Enabling the multicast protocol(s) on configured interfaces.

For PIM for IPv6 to operate, the Multicast Listener Discovery (MLD) protocol must be enabled. You must also configure a unicast routing protocol, such as OSPFv3. For both DVMRP and PIM for IPv4 to operate, IGMP must be enabled. See “[Configuring IGMP](#)” on page 23-20 for IGMP configuration information.

## Understanding MLD

Multicast allows a source to send a single copy of data using a single IP address from a well-defined range for an entire group of recipients (an MLD group). A source sends data to an MLD group by simply setting the destination IP address of the datagram to be the MLD group address. Sources do not need to register in any way before they can begin sending data to a group, and do not need to be members of the group themselves. Routers between the source and recipients use the group address to route the data, forwarding duplicate data packets only when the path to recipients diverges.

Hosts that wish to receive IPv6 data from the MLD group join the group by sending a message to a multicast router on a local interface, using MLD.

Multicast routers communicate among themselves using a multicast routing protocol, such as DVMRP or PIM-SM. These protocols calculate a multicast distribution tree of recipients to ensure that:


- Multicast traffic reaches all recipients that have joined the MLD group
- Multicast traffic does not reach networks that do not have any such recipients (unless the network is a transit network on the way to other recipients)
- The number of identical copies of the same data flowing over the same link is minimized.

Group membership management is fundamental to the multicasting process. An arbitrary group of receivers can express interest in receiving a particular multicast stream, regardless of the physical or geographical boundaries of its members.

The purpose of IPv6 multicast group management is to optimize a switched network’s performance so multicast packets will only be forwarded to those ports containing MLD group hosts or multicast switch devices instead of flooding to all ports in the subnet (VLAN).

MLD uses three key components to control multicast membership:

- **Source** — A server that sends an IPv6 multicast data stream with a particular MLD destination IPv6 and MAC address. A server may not have direct MLD involvement, as it often does not receive a multicast stream, but only sends a multicast stream.
- **Querier** — A device that periodically sends out queries in search of multicast hosts on a directly connected network. If multiple queriers are present on the LAN, the querier with the lowest IP address assumes the role.
- **Host** — A client end station that sends one of two MLD messages to a querier:
  - Join message — Indicates the host wants to receive transmissions associated to a particular multicast group.
  - Leave message — Indicates the host wants to stop receiving the multicast transmissions.

**Figure 24-1 MLD Querier Determining Group Membership**

As shown in [Figure 24-1](#), an MLD-enabled device can periodically ask its hosts if they want to receive multicast traffic. If there is more than one device on the LAN performing IPv6 multicasting, one of these devices is elected querier and assumes the responsibility of querying the LAN for group members.

Based on the group membership information learned from MLD, a device can determine which (if any) multicast traffic needs to be forwarded to each of its ports. At Layer 3, multicast switch devices use this information, along with a multicast routing protocol, to support IPv6 multicasting across the Internet.

MLD provides the final step in IPv6 multicast delivery. It is only concerned with forwarding multicast traffic from the local switch device to group members on a directly attached subnetwork or LAN segment.

MLD neither alters nor routes any IPv6 multicast packets. Since MLD is not concerned with the delivery of IPv6 multicast packets across subnetworks, an external IPv6 multicast device is needed if IPv6 multicast packets have to be routed across different subnetworks.

## MLD Support on Extreme Networks Devices

Extreme Networks devices implement MLD versions 1 and 2, which includes interoperability with version 1 hosts.

Depending on your Extreme Networks device, MLD can be configured independently at the switch level (Layer 2) and at the router level (Layer 3).

Extreme Networks devices support MLD as follows:

- Passively snooping on the MLD query and MLD report packets transferred between IPv6 multicast switches and IPv6 multicast host groups to learn MLD group members. Each Layer 2 device records which ports MLD packets are received on, depending on the kind of MLD message, so multicast data traffic is not flooded across every port on the VLAN when it is received by the switch.

MLD snooping is disabled by default on Extreme Networks devices. You can enable it using the `set mld enable` command on Extreme Networks 7100-Series devices as described in “Configuring MLD” on page 24-5.

- MLD actively sends MLD query messages to learn locations of multicast switches and member hosts in MLD groups within each VLAN.

## Example: Sending a Multicast Stream

**Figure 24-2 Sending a Multicast Stream with No Directly Attached Hosts**

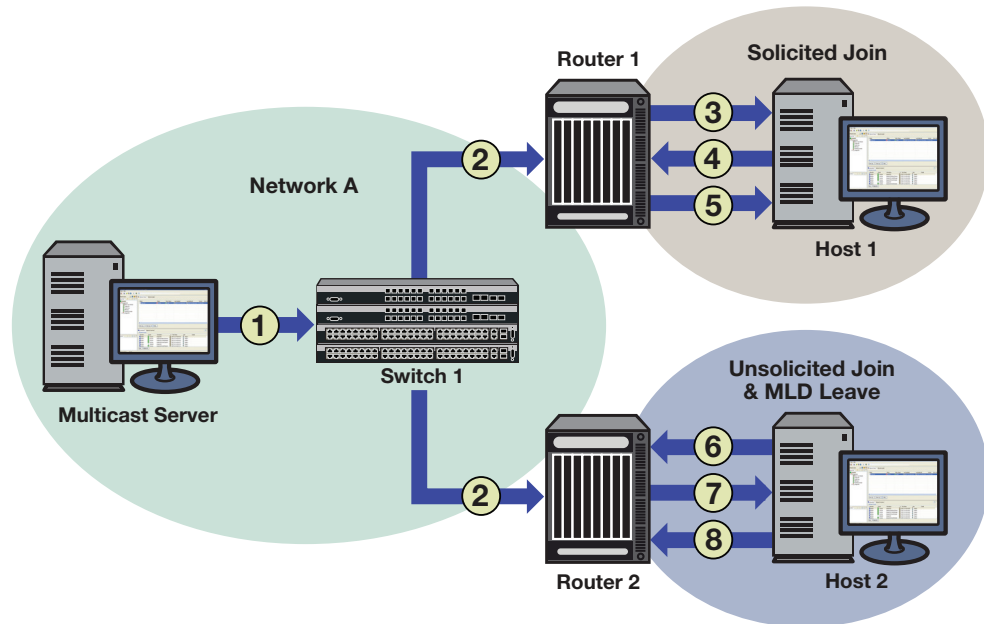


Figure 24-2 provides an example of MLD processing on Extreme Networks devices when there are no directly attached hosts.

1. A single IP multicast server, with no directly attached hosts, sends a multicast stream into the network via Switch 1.
2. Because MLD snooping is disabled, Switch 1 floods the multicast stream to all ports which are linked to Router 1 and Router 2.

Each router performs an MLD forwarding check to see if there are any hosts that want to join the multicast group on its locally attached network. Each router drops multicast packets until a host joins the group using one of the following messages:

- **solicited join** (sent in response to an MLD query produced by the router’s interface)

In Figure 24-2, this type of exchange occurs between Router 1 and Host 1 when:

- (3) Router 1 sends a query to potential Host 1.
- (4) Host 1 responds with a join message.
- (5) Router 1 forwards the multicast stream.

- **unsolicited join** (sent as a request without receiving an MLD query first)

In Figure 24-2, this type of exchange occurs between Router 2 and Host 2 when:

- (6) Host 2 sends a join message to Router 2.

- (7) Router 2 forwards the multicast stream to Host 2.
- (8) When it no longer wants to receive the stream, Host 2 can do one of the following:
  - Send a leave message to Router 2.
  - Time out the MLD entry by not responding to further queries from Router 2.

## Configuring MLD

MLD is configured in switch mode on Extreme Networks 7100-Series devices. At Layer 2, MLD can be enabled for VLANs, regardless of whether it is enabled on routed interfaces. If, however, MLD is enabled on a routed interface, and the routed interface is a routed VLAN, then MLD must also be enabled at the switch level.

### MLD Configuration Commands

Table 24-1 lists the MLD configuration commands for Extreme Networks 7100-Series devices.

**Table 24-1 MLD Configuration Commands**

Task	Command
Enable MLD on one or more VLANs.	<b>set mld enable</b> <i>vlan-list</i>
Disable MLD on one or more VLANs.	<b>set mld disable</b> <i>vlan-list</i>
Enable MLD querying on one or more VLANs.	<b>set mld query-enable</b> <i>vlan-list</i>
Disable MLD querying on one or more VLANs.	<b>set mld query-disable</b> <i>vlan-list</i>
Determine what action to take with multicast frames when the MLD flow table is full.	<b>set mld flow-full-action</b> <i>action</i>
Configure MLD settings on one or more VLANs.	<b>set mld config</b> <i>vlan-list</i> {[ <b>query-interval</b> <i>query-interval</i> ] [ <b>mld-version</b> <i>mldversion</i> ] [ <b>max-resp-time</b> <i>max-resp-time</i> ] [ <b>robustness</b> <i>robustness</i> ] [ <b>last-mem-int</b> <i>last-mem-int</i> ] [ <b>fast-leave</b> <i>fast-leave</i> ] [ <b>filter-id</b> <i>filter-id</i> ] [ <b>filter-status</b> {enable   disable}]}
Remove MLD configuration settings for one or more VLANs.	<b>set mld delete</b> <i>vlan-list</i>
Create a new static MLD entry or add one or more new ports to an existing entry.	<b>set mld static</b> <i>group vlan-list</i> [ <b>modify</b> {[ <b>include-ports</b> ] [ <b>exclude-ports</b> ]}]
Remove a static MLD entry or remove one or more ports from an existing entry.	<b>clear mld static</b> <i>group vlan-list</i> [ <b>modify</b> {[ <b>include-ports</b> ] [ <b>exclude-ports</b> ]}]
Change the MLD classification of received IP frames.	<b>set mld protocols classification</b> <i>classification protocol-id protocol-id</i> [ <b>modify</b> ]
Clear the binding of IP protocol ID to MLD classification.	<b>clear mld protocols protocol-id</b> <i>protocol-id</i>
Enable port fast leave on the specified port or range of ports.	<b>set mld portFastLeave</b> <i>port-list</i>
Disable port fast leave on the specified port or range of ports.	<b>clear mld portFastLeave</b> <i>port-list</i>
Create an input filter to apply to the VLAN.	<b>set mld input-filter</b> <i>filter-id rule-id start-ip ip-address end-ip ip-address protocol-action</i> {deny   allow} <b>flow-action</b> {drop   flood   allow}

**Table 24-1 MLD Configuration Commands (continued)**

Task	Command
Clear an input filter.	<b>clear mld input-filter</b> <i>filter-id</i> [ <i>rule-id</i> ]
Set the action taken when the first few frames of a multicast stream are received (that is, before the stream is added to the MLD database).	<b>set mld unknown-input-action</b> {routers   flood   discard}

## Basic MLD Configurations

[Procedure 24-1](#) describes the basic steps to configure MLD on Extreme Networks 7100-Series devices. This procedure assumes that the VLANs on which MLD will run have been configured and enabled with IP interfaces.

**Procedure 24-1 Basic MLD Configuration**

Step	Task	Command
1.	In switch mode, configure MLD for each VLAN interface.	<b>set mld config</b> <i>vlan-list</i> {[ <b>query-interval</b> <i>query-interval</i> ] [ <b>mld-version</b> <i>mldversion</i> ] [ <b>max-resp-time</b> <i>max-resp-time</i> ] [ <b>robustness</b> <i>robustness</i> ] [ <b>last-mem-int</b> <i>last-mem-int</i> ]}
2.	In switch mode, enable MLD on each VLAN interface.	<b>set mld enable</b> <i>vlan-list</i>
3.	In switch mode, enable MLD querying on each of the VLANs specified in step 2.	<b>set mld query-enable</b> <i>vlan-list</i>

For more information on MLD CLI commands, refer to your device's *CLI Reference Guide*.

## Example MLD Configuration

The following example enables MLD and MLD query on VLANs 2 and 3:

```
7100 Chassis->set mld enable 2,3
7100 Chassis->set mld query-enable 2,3
```

## MLD Display Commands

[Table 24-2](#) lists MLD show commands.

**Table 24-2 MLD Show Commands**

Task	Command
Display the status of MLD on one or more VLANs.	<b>show mld enable</b> <i>vlan-list</i>
Display the MLD query status of one or more VLANs.	<b>show mld query</b> <i>vlan-list</i>
Display the action to be taken with multicast frames when the multicast MLD flow table is full.	<b>show mld flow-full-action</b>
Display MLD configuration information for one or more VLANs.	<b>show mld config</b> <i>vlan-list</i>
Display MLD information regarding multicast group membership.	<b>show mld groups</b> [ <b>group</b> <i>group</i> ] [ <b>vlan-list</b> <i>vlan-list</i> ] [ <b>sip</b> <i>sip</i> ] [ <b>-verbose</b> ]

**Table 24-2 MLD Show Commands (continued)**

<b>Task</b>	<b>Command</b>
Display static MLD ports for one or more VLANs or MLD groups.	<b>show mld static</b> [group <i>group</i> ] [vlan-list <i>vlan-list</i> ]
Display the binding of IP protocol ID to MLD classification.	<b>show mld protocols</b>
Display MLD information for a specific VLAN.	<b>show mld vlan</b> [ <i>vlan-list</i> ]
Display MLD reporter information.	<b>show mld reporters</b> [portlist <i>portlist</i> ] [group <i>group</i> ] [vlan-list <i>vlan-list</i> ] [sip <i>sip</i> ]
Display MLD flow information.	<b>show mld flows</b> [portlist <i>portlist</i> ] [group <i>group</i> ] [vlan-list <i>vlan-list</i> ] [sip <i>sip</i> ]
Display MLD counter information.	<b>show mld counters</b>
Display the number of MLD flows set on the Extreme Networks 7100-Series device.	<b>show mld number-flows</b>
Display configuration information for input filters.	<b>show mld input-filter</b> [ <i>filter-id</i> ] [ <i>rule-id</i> ]
Display the action taken when the first frames of a multicast stream are received.	<b>show mld unknown-input-action</b>





---

## System Logging Configuration

This chapter provides the following information about configuring and monitoring Syslog on Extreme Networks 7100-Series devices.

For information about...	Refer to page...
<a href="#">Using Syslog in Your Network</a>	25-1
<a href="#">Syslog Overview</a>	25-2
<a href="#">Configuring Syslog</a>	25-6

### Using Syslog in Your Network

Syslog, short for System Logging, is a standard for forwarding log messages in an IP network that is typically used for network system management and security auditing. The term often applies to both the actual Syslog protocol, as well as the application sending Syslog messages.

As defined in RFC 3164, the Syslog protocol is a client/server-type protocol which enables a station or device to generate and send a small textual message (less than 1024 bytes) to a remote receiver called the Syslog server. Messages are transmitted using User Datagram Protocol (UDP) packets and are received on UDP port 514. These messages inform about simple changes in operational status or warn of more severe issues that may affect system operations.

When managed properly, logs are the eyes and ears of your network. They capture events and show you when problems arise, giving you information you need to make critical decisions whether you are building a policy rule set, fine tuning an Intrusion Detection System, or validating which ports should be open on a server. However, since it's practically impossible to wade through the volumes of log data produced by all your servers and network devices, Syslog's ability to place all events into a single format so they can be analyzed and correlated makes it a vital management tool. Because it is supported by a wide variety of devices and receivers across multiple platforms, you can use it to integrate log data from many different types of systems into a central repository.

Efficient Syslog monitoring and analysis reduces system downtime, increases network performance, and helps tighten security policies. It can help you:

- Troubleshoot switches, firewalls and other devices during installation and problem situations.
- Perform intrusion detection.
- Track user activity.

## Syslog On 7100-Series Switches

By default, Syslog is operational on 7100-Series devices at startup. All generated messages are eligible for logging to local destinations and to remote servers configured as Syslog servers. Using simple CLI commands, you can adjust device defaults to configure the following:

- Message sources— which system applications on which modules should log messages
- Message destinations— will messages be sent to the local console, the local file system, or to remote Syslog servers? Which facility (functional process) will be allowed to send to each destination?

The following section provides an overview of Syslog features and functions supported on Extreme Networks devices and their default configurations. Later sections will provide instructions on changing default settings to suit your network logging needs.

## Syslog Overview

Developers of various operating systems, processes, and applications determine the circumstances that will generate system messages and write those specifications into their programs. Messages can be generated to give status, either at a certain period of time, or at some other interval, such as the invocation or exit of a program. Messages can also be generated due to a set of conditions being met. Typically, developers quantify these messages into one of several broad categories, generally consisting of the facility that generated them, along with an indication of the severity of the message. This allows system administrators to selectively filter the messages and be presented with the more important and time sensitive notifications quickly, while also having the ability to place status or informative messages in a file for later review.

Switches must be configured with rules for displaying and/or forwarding event messages generated by their applications. In addition, Syslog servers need to be configured with appropriate rules to collect messages so they can be stored for future reference. This document will describe how to complete these key configuration steps on 7100-Series platforms.

If C2 security mode is enabled, while in Read-Write user mode you can not:

- Create, modify, or clear a server logging configuration
- Create, modify, or clear a default logging configuration
- Create, modify, or clear a logging application configuration

## Configuring Syslog Message Disposition

The Syslog implementation on 7100-Series devices uses a series of system logging messages to track device activity and status. These messages inform users about simple changes in operational status or warn of more severe issues that may affect system operations. Logging can be configured to display messages at a variety of different severity levels about application-related error conditions occurring on the device.

You can decide to have all messages stored locally, as well as to have all messages of a high severity forwarded to another device. You can also have messages from a particular facility sent to some or all of the users of the device, and displayed on the system console. For example, you may want all messages that are generated by the mail facility to be forwarded to one particular Syslog server. However you decide to configure the disposition of the event messages, the process of having them sent to a Syslog collector generally consists of:

- Determining which messages at which severity levels will be forwarded.
- Defining one or more remote receivers (Syslog servers/console displays).

## Filtering by Severity and Facility

Syslog daemons determine message priority by filtering them based on a combined facility and severity code. Severity indicates the seriousness of the error condition generating the Syslog message. This is a value from 1 to 8, with 1 indicating highest severity. Facility categorizes which functional process is generating an error message. The Extreme Networks implementation uses the eight facility designations reserved for local use: **local0** – **local7** defined in RFC 3164. You can modify these default facility and severity values to control message receipt and aid in message sorting on target servers.

For example, you can configure all router messages to go to Server 1 using facility local1, while all SNMP messages go to Server 1 using facility local2.

The following sections provide greater detail on modifying key Syslog components to suit your enterprise.

## Syslog Components and Their Use

Table 25-1 describes the Extreme Networks implementation of key Syslog components.

**Table 25-1 Syslog Terms and Definitions**

Term	Definition	Extreme Networks Usage
Facility	Categorizes which functional process is generating an error message. Syslog combines this value and the severity value to determine message priority.	Extreme Networks uses the eight facility designations reserved for local use: <b>local0</b> – <b>local7</b> . Default is <b>local4</b> , which allows the message severity portion of the priority code to be visible in clear text, making message interpretation easiest. For more information about facility designations, refer to RFC 3164.
Severity	Indicates the severity of the error condition generating the Syslog message. The lower the number value, the higher will be the severity of the condition generating the message.	Extreme Networks devices provide the following eight levels: <ul style="list-style-type: none"> <li>1 - emergencies (system is unusable)</li> <li>2 - alerts (immediate action required)</li> <li>3 - critical conditions</li> <li>4 - error conditions</li> <li>5 - warning conditions</li> <li>6 - notifications (significant conditions)</li> <li>7 - informational messages</li> <li>8 - debugging messages</li> </ul> The default Syslog configuration allows applications (log message sources) to forward messages at a severity level of 6, and destinations (console, file system, or remote Syslog servers) to log messages at a severity level of 8.



**Note:** Numerical values used in Extreme Networks syslog CLI and the feature's configuration MIB range from 1-8. These map to the RFC 3164 levels of 0-7 respectively. Syslog messages generated report the RFC 3164 specified level values.

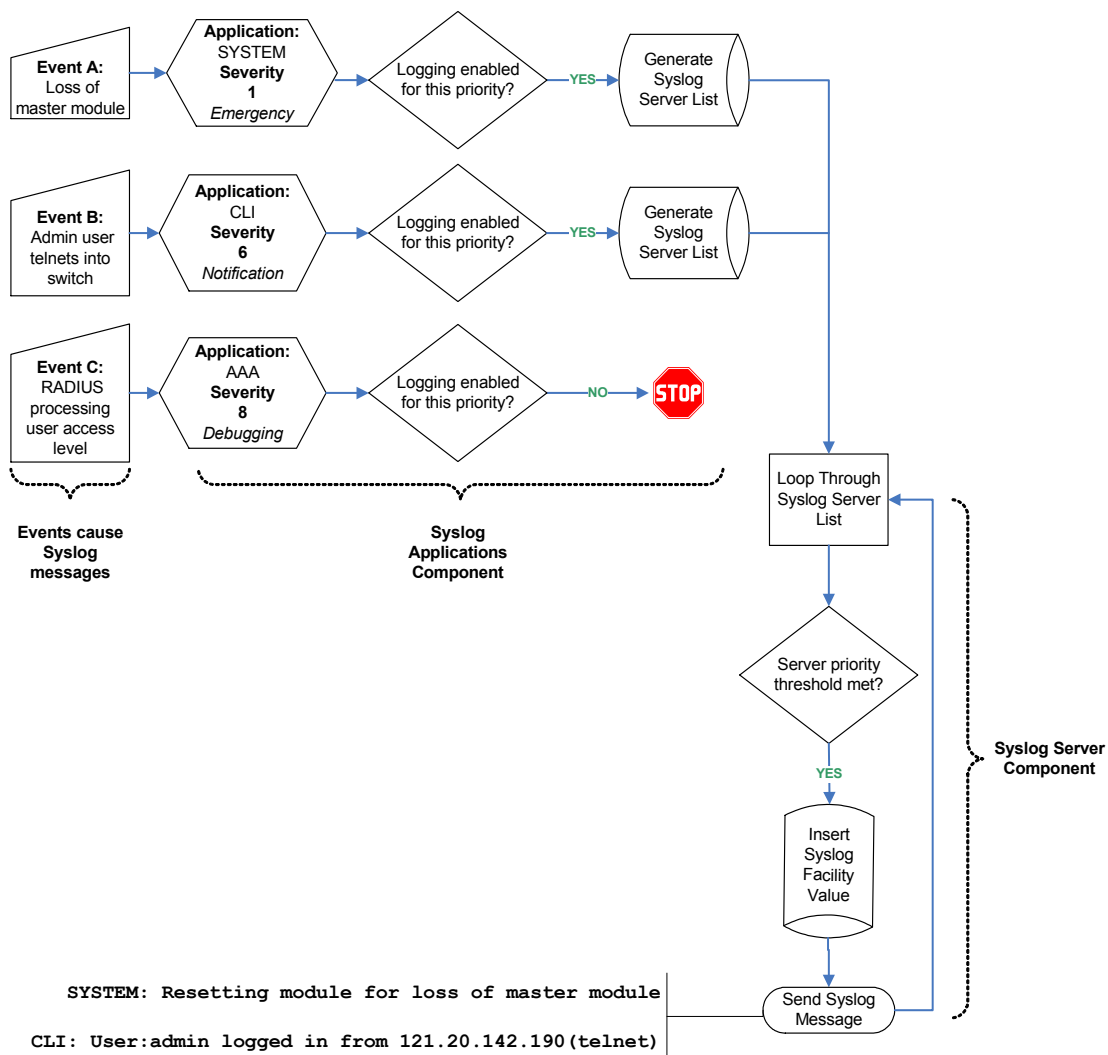
**Table 25-1 Syslog Terms and Definitions (continued)**

Term	Definition	Extreme Networks Usage
Application	Client software applications running on devices that can generate Syslog messages.	<p>Extreme Networks supported applications and their associated CLI mnemonic values include:</p> <p><b>CLI</b> - Command Line Interface</p> <p><b>SNMP</b> - Simple Network Management Protocol</p> <p><b>Webview</b> - Extreme Networks Web-based system management</p> <p><b>System</b> - System messages</p> <p><b>RtrFe</b> - Router Forwarding Engine</p> <p><b>Trace</b> - Trace logging</p> <p><b>RtrLSNat</b> - Load Share Network Address Translation</p> <p><b>FlowLimt</b> - Flow limiting</p> <p><b>UPN</b> - User Personalized Networks</p> <p><b>AAA</b> - Authentication, Authorization and Accounting</p> <p>Use the <b>show logging application all</b> command to list supported applications and the corresponding CLI numeric or mnemonic values you can use to configure application logging on your devices.</p>
Syslog server	A remote server configured to collect and store Syslog messages.	<p>Extreme Networks devices allow up to 8 server IP addresses to be configured as destinations for Syslog messages. By default, Syslog server is globally enabled, with no IP addresses configured, at a severity level of 8.</p>

## Basic Syslog Scenario

Figure 25-1 on page 25-5 shows a basic scenario of how Syslog components operate on an Extreme Networks switch. By default, all applications running on the Extreme Networks switch are allowed to forward Syslog messages generated at severity levels 6 through 1. In the configuration shown, these default settings have not been changed.

Figure 25-1 Basic Syslog Scenario



Default application settings in the example in [Figure 25-1](#) have not been modified. Therefore, an emergency message triggered by a system reset due to loss of the master module is forwarded to Syslog destinations. The CLI-related message notifying that a user has logged in remotely is also forwarded. Configured Syslog server(s) will receive all forwarded messages since their default severity threshold is at 8 (accepting messages at all severity levels).

Any messages generated by applications at severity levels 7 and 8 are not forwarded in this example. For instance, forwarding does not occur for an AAA authentication-related debugging message with information about RADIUS access level processing for a particular user. If at some point in time it becomes necessary, for example, to log all AAA authentication-related message activity and to save it to a file so authentication details can be tracked, the administrator can allow that specific application to forward debugging messages to a Syslog server, as well as to the console and persistent file storage.

For more information on how to configure these basic settings, refer to [“Syslog Command Precedence”](#) on page 25-7, and the [“Configuration Examples”](#) on page 25-12.

## Interpreting Messages

Every system message generated by the 7100-Series platforms follows the same basic format:

```
<facility/severity> time stamp address application [slot] message text
```

### Example

This example shows Syslog informational messages, displayed with the **show logging buffer** command. It indicates that messages were generated by facility code 16 (local4) at severity level 5 from the CLI application on IP address 10.42.71.13.

```
7100 Chassis(rw)->show logging buffer
<165>Sep  4 07:43:09 10.42.71.13 CLI[5]User:rw logged in from 10.2.1.122 (telnet)
<165>Sep  4 07:43:24 10.42.71.13 CLI[5]User: debug failed login from 10.4.1.100
(telnet)
```

[Table 25-2](#) describes the components of these messages.

**Table 25-2 Syslog Message Components**

Component	Description	Example Code
Facility/Severity	Combined code indicating the facility generating the message and the severity level used to determine message priority. Facility codes 16 - 23 are Syslog designations for local0 - local7, the Extreme Networks supported designations for local use. For a complete list of facility codes, refer to RFC 3164.	<165> = Numerical code indicating a message from facility local4 at severity 5.
Time stamp	Month, date, and time the Syslog message appeared.	Sep 4 07:43:09
Address	IP address of the client originating the Syslog message.	10.42.71.13
Application	Client process generating the Syslog message.	CLI
Slot/Module	Slot location of the device module generating the Syslog message.	(5) = Slot 5 in the chassis.
Message text	Brief description of error condition.	User: debug failed login from 10.4.1.100 (telnet)

## Configuring Syslog

For information about...	Refer to page...
<a href="#">Syslog Command Precedence</a>	<a href="#">25-7</a>
<a href="#">About Server and Application Severity Levels</a>	<a href="#">25-7</a>
<a href="#">Configuring Syslog Server(s)</a>	<a href="#">25-7</a>
<a href="#">Modifying Syslog Server Defaults</a>	<a href="#">25-8</a>
<a href="#">Reviewing and Configuring Logging for Applications</a>	<a href="#">25-9</a>
<a href="#">Enabling Console Logging and File Storage</a>	<a href="#">25-10</a>
<a href="#">Configuration Examples</a>	<a href="#">25-12</a>

## Syslog Command Precedence

Table 25-3 lists basic Syslog commands and their order of precedence on Extreme Networks switches.

**Table 25-3 Syslog Command Precedence**

Syslog Component	Command	Function
Logging defaults	<b>set logging default</b> { [ <b>facility</b> <i>facility</i> ] [ <b>severity</b> <i>severity</i> ] [ <b>port</b> <i>port</i> ] }	Sets default parameters for facility code, severity level and/or UDP port for all Syslog servers and local destinations.  Settings will be applied when Syslog servers are configured without specifying values with the <b>set logging server</b> command. This command overrides factory defaults.
Server settings	<b>set logging server</b> <i>index</i> <b>ip-addr</b> <i>ip-addr</i> [ <b>facility</b> <i>facility</i> ] [ <b>severity</b> <i>severity</i> ] [ <b>descr</b> <i>descr</i> ] [ <b>port</b> <i>port</i> ] <b>state</b> <b>enable</b>   <b>disable</b>	During or after new server setup, specifies a server index, IP address, and operational state for a Syslog server. Optionally, this command specifies a facility code, severity level at which messages will be accepted, text string description, and/or UDP port for the specified server.  This command overrides system defaults for the specified server. If not specified with this or the <b>set logging default</b> command, optional server parameters will be set to the system defaults listed in Table 25-4 on page 8.
Application settings	<b>set logging application</b> {[ <i>mnemonic</i>  all]} [ <b>level</b> <i>level</i> ] [ <b>servers</b> <i>servers</i> ]	Sets the severity level at which one or all applications will send messages to Syslog servers. If not specified, settings will apply to all configured servers and severity level will not be changed from system defaults.

## About Server and Application Severity Levels

By default, client applications will forward Syslog messages at severity levels 6 through 1, and servers will log messages at all severity levels (8 through 1). You can use the procedures described in this chapter to change these parameters, fine tuning the scope of message logging and modifying the Syslog behavior between one or more client applications and one or more servers.

## Configuring Syslog Server(s)

Use the following commands to configure one or more servers as destinations for Syslog messages and verify the configuration:

1. Add a Syslog server to the device's server list:

```
set logging server index ip-addr ip-addr state enable
```

*Index* is a value from 1 to 8 that specifies the server table index number for this server.

2. (Optional) Verify the server configuration:

```
show logging server [index]
```

If *index* is not specified, information for all configured Syslog servers will be displayed.



## Example

This sample output from the **show logging server** command shows that two servers have been added to the device's Syslog server list. These servers are using the default UDP port 514 to receive messages from clients and are configured to log messages from the local1 and local2 facilities, respectively. Logging severity on both servers is set at 5 (accepting messages at severity levels 5 through 1). Using the commands described in the next section, these settings can be changed on a per-server basis, or for all servers.

```
7100 Chassis(rw)->show logging server
```

	IP Address	Facility	Severity	Description	Port	Status
1	132.140.82.111	local1	warning(5)	default	514	enabled
2	132.140.90.84	local2	warning(5)	default	514	enabled

## Modifying Syslog Server Defaults

Unless otherwise specified, the switch will use the default server settings listed in [Table 25-4](#) for its configured Syslog servers:

**Table 25-4 Syslog Server Default Settings**

Parameter	Default Setting
facility	local4
severity	8 (accepting all levels)
descr	no description applied
port	UDP port 514

Use the following commands to change these settings either during or after enabling a new server.

## Displaying System Logging Defaults

To display system logging defaults, or all logging information, including defaults:

```
show logging {default|all}
```

## Modifying Default Settings

You can change factory default logging settings using one of the following methods.

- To specify logging parameters during or after new server setup:

```
set logging server index ip-addr ip-addr [facility facility] [severity severity] [descr descr] [port port] state enable
```

If not specified, optional server parameters will be set to the system defaults listed in [Table 25-4](#). Refer back to [Filtering by Severity and Facility](#) and to [Table 25-1](#) for more information on how these parameters operate.

- To change default parameters for all servers:

```
set logging default {[facility facility] [severity severity] [port port]}
```

## Examples

This example shows how to configure the switch to forward messages from facility category local6 at severity levels 3, 2, and 1 to Syslog server 1 at IP address 134.141.89.113:

```
7100 Chassis(rw)->set logging server 1 ip-addr 134.141.89.113 facility local6
severity 3
```

This example shows how to change Syslog defaults so that messages from the local2 facility category at a severity level of 4 will be forwarded to all servers. These settings will apply to all newly-configured servers, unless explicitly configured with the **set logging server** command:

```
7100 Chassis(rw)->set logging default facility local2 severity 4
```

## Reviewing and Configuring Logging for Applications

By default, all applications running on 7100-Series devices are allowed to forward messages at severity levels 6 through 1 to all configured destinations (Syslog servers, the console, or the file system).

### Displaying Current Application Severity Levels

To display logging severity levels for one or all applications currently running on your device:

```
show logging application {mnemonic|all}
```

#### Example

This example shows output from the **show logging application all** command. A numeric and mnemonic value for each application is listed with the severity level at which logging has been configured and the server(s) to which messages will be sent. In this case, logging for applications has not been changed from the default severity level of 6. This means that notifications and messages with severity values 6 through 1 will be sent to configured servers.

```
7100 Chassis(rw)->show logging application all
```

Application	Current Severity Level	Server List	
88	RtrAcl	6	1-8
89	CLI	6	1-8
90	SNMP	6	1-8
91	Webview	6	1-8
93	System	6	1-8
95	RtrFe	6	1-8
96	Trace	6	1-8
105	RtrLSNat	6	1-8
111	FlowLimt	6	1-8
112	UPN	6	1-8
117	AAA	6	1-8
118	Router	6	1-8
140	AddrNtfy	6	1-8
141	OSPF	6	1-8
142	VRRP	6	1-8
145	RtrArpProc	6	1-8
147	LACP	6	1-8

148	RtrNat	6	1-8
151	RtrTwcb	6	1-8
158	HostDoS	6	1-8

1 (emergencies)    2 (alerts)    3 (critical)  
 4 (errors)        5 (warnings)    6 (notifications)  
 7 (information)    8 (debugging)



**Note:** Mnemonic values are case sensitive and must be typed as they are listed in the **show logging application** command display for your device. Refer to [Table 25-1](#) for sample CLI mnemonic values.

## Modifying Severity Levels and Assigning Syslog Servers for Applications

Applications running on Extreme Networks devices will use the default Syslog settings unless otherwise configured by the **set logging server** or **set logging default** commands as previously described.

To modify the severity level at which log messages will be forwarded and the server(s) to which they will be sent for one or all applications:

```
set logging application {[mnemonic|all]} [level level] [servers servers]
```

### Example

This example shows how to set the severity level for SSH (Secure Shell) to 5 so that warning conditions and messages of greater severity (levels 5 to 1) generated by that application will be sent to Syslog server 1.

```
7100 Chassis(rw)->set logging application SSH level 5 server 1
```

## Enabling Console Logging and File Storage

7100-Series devices allow you to display logging messages to the console and save to a persistent file. In addition, 7100-Series devices also provide the option of allowing you to display messages to the current console CLI session only.

Console logging allows you to view only as many messages as will fit on the screen. As new messages appear, old messages simply scroll off the console. While this is a temporary means of logging information, it allows you to track very specific activities quickly and easily. Console log messages can also be saved to a persistent file at two locations:

- slotX/logs/current.log — Location of current system log messages (up to 256k), where X specifies the slot location of the device.
- slotX/logs/old.log — Location of previous system log messages, where X specifies the slot location of the device. Current messages will be moved to the old.log when current.log file exceeds 256k.

Use the following commands to review and configure console logging and file storage.

### Displaying to the Console and Saving to a File

To display log messages to the console and save to a persistent file:

```
set logging local console {enable | disable} file {enable | disable} sfile {enable | disable}
```



**Note:** The **set logging local** command requires that you specify both console and file settings. For example, **set logging local console enable** would not execute without also specifying **file enable** or **disable** or **sfile enable** or **disable** for a secure file.

## Displaying to the Current CLI Session

To display logging to the current CLI console session on a 7100-Series device:

```
set logging here enable
```

This adds the current CLI session to the list of Syslog destinations, and will be temporary if the current CLI session is using Telnet or SSH.

## Displaying a Log File

To display the contents of the persistent log file:

```
show file slotslotnumber/logs/current.log|old.log
```



**Note:** These log files may also be copied to another device using FTP or TFTP.

## CLI and SNMP Audit Logging

CLI and SNMP audit logging provide Syslog messages for:

- CLI configuration changes
- SNMP configuration changes
- SNMP authentication failures



**Note:** In Release 7.41, CLI and SNMP audit logging messages were moved from other Syslog applications to the security application.

CLI and SNMP audit logging is turned on by configuring logging for the security application at level 7 (information), using the **set logging application security level 7** command. CLI audit logging messages display the user mode, the IP address of the station originating the CLI command and the type of client connection, the CLI command entered and a status indicating command success (OK) or failure (Fail). The logging output when entering this command would log a message similar to:

```
<166>May 3 16:44:12 10.21.130.166 Security[1]User:admin;
Source:134.141.90.54(ssh); Action:"set logging application security level 7 ";
Status:OK
```

SNMP audit logging messages display the SNMP user, the IP address of the station originating the SET, the MIB leaf modified and the setting value, and a status indicating success (OK) or failure (Fail).

SNMP authentication failure messages display the SNMP user, the IP address of the station originating the request, and the failure reason.

## Configuration Examples

### Enabling a Server and Console Logging

[Procedure 25-1](#) shows how you would complete a basic Syslog configuration. In this example, the default application severity level has not been modified, allowing all applications to forward messages to configured destinations. One Syslog server is configured on IP address 10.1.1.2, logging all messages. Console logging is enabled, but persistent file storage is not.

#### Procedure 25-1 Configuring a Server and Console Logging

Step	Task	Command(s)
1.	Configure Syslog server 1 and accept default settings (listed in <a href="#">Table 25-4</a> on page 25-8).	<b>set logging server 1 ip-addr 10.1.1.2 state enable</b>
2.	(Optional) Verify that application logging settings are at default values for the enabled server.	<b>show logging application all</b>
3.	Enable console logging and disable file storage.	<b>set logging local console enable file disable</b>



**Note:** The **set logging local** command requires that you specify both console and file settings. For example, **set logging local console enable** would not execute without also specifying **file enable** or **disable**.

### Adjusting Settings to Allow for Logging at the Debug Level

[Procedure 25-2](#) shows how you would adjust the previous Syslog configuration so that all AAA-related authentication messages (level 8) could be forwarded to Server 2 at IP address 10.1.1.3, displayed on the console and saved to persistent file storage. This would enable all Syslog messaging capabilities for this particular application. Since the severity for this new server has not changed from the default of level 8, there is no need to adjust this setting.

#### Procedure 25-2 Adjusting Settings for an Application

Step	Task	Command(s)
1.	Configure Syslog server 2 and accept default settings (listed in <a href="#">Table 25-4</a> on page 25-8).	<b>set logging server 2 ip-addr 10.1.1.3 state enable</b>
2.	Set the severity level for the AAA application to level 8.	<b>set logging application AAA level 8 servers 2</b>
3.	Enable console logging and file storage.	<b>set logging local console enable file enable</b>

---

## Network Monitoring Configuration

This document describes the network monitoring features and their configuration on Extreme Networks 7100-Series devices.

For information about...	Refer to page...
<a href="#">Using Network Monitoring in Your Network</a>	26-1
<a href="#">Network Monitoring Overview</a>	26-2
<a href="#">Configuring Network Monitoring</a>	26-7

### Using Network Monitoring in Your Network

7100-Series network monitoring features support for:

- Console/Telnet based monitoring:
  - Display contents and determine the size of the command history buffer
  - Close an active console port or Telnet session
- Network Diagnostics:
  - Determine the availability of another node on the network (ping)
  - Display a hop-by-hop path through an IP network (traceroute)
  - Query name servers (nslookup)
- Display of switch network connections:
  - Display statistics for the active connections on the switch
  - Display switch users
  - Send a message to switch user
- SMON statistics:
  - Monitor SMON priority and VLAN statistics counting
- RMON:
  - Record statistics measured by the RMON probe for each monitored interface on the device.
  - Record periodic statistical samples from a network.
  - Periodically gather statistical samples from variables in the probe and compares them with previously configured thresholds.
  - Record statistics associated with each host discovered on the network.

- Control the generation and notification of events from the device.
- Generate tables that describe hosts that top a list ordered by one of their statistics.
- Record statistics for conversations between two IP addresses.
- Allow packets to be matched by a filter equation.
- Allow packets to be captured upon a filter match.

## Network Monitoring Overview

This section provides an overview of network monitoring configuration.

### Console/Telnet History Buffer

The history buffer lets you recall your previous CLI input. The size of the history buffer determines how many lines of previous CLI input are available for recall. By default, the size of this buffer is 20 lines. The configured size can be displayed. The contents of the buffer can be displayed.

Use the **set history** command in any command mode to set the size of the history buffer to a value between **1 - 100** lines.

```
7100 Chassis(rw)->set history 25
```

Use the **show history** command in any command mode to display the currently configured size of the history buffer.

```
7100 Chassis(rw)->show history
History size currently set to: 25
7100 Chassis(rw)->
```

Use the **history** command in any command mode to display the contents of the history buffer.

```
7100 Chassis(rw)->history
 1 history
 2 show gvrp
 3 show vlan
 4 show igmp
7100 Chassis(rw)->
```

### Network Diagnostics

7100-Series network diagnostics provide for:

- Pinging another node on the network to determine its availability
- Performing a traceroute through the IP network to display a hop-by-hop path from the device to a specific destination host
- Querying name servers to translate hostnames to IP addresses or IP addresses to hostnames

Use the **ping** command, in any command mode, to determine whether the specified node is available.

```
7100 Chassis(rw)->ping -c 10 127.0.0.1
PING 127.0.0.1 (localhost) 64 bytes of data.
64 bytes from 127.0.0.1 (localhost): icmp_seq=0 ttl=64 time=1.58 ms
```

```
64 bytes from 127.0.0.1 (localhost): icmp_seq=1 ttl=64 time=1.52 ms
64 bytes from 127.0.0.1 (localhost): icmp_seq=2 ttl=64 time=1.57 ms
64 bytes from 127.0.0.1 (localhost): icmp_seq=3 ttl=64 time=2.26 ms
64 bytes from 127.0.0.1 (localhost): icmp_seq=4 ttl=64 time=1.42 ms
```

Use the **tracert** command, in any command mode, to display a hop-by-hop path through an IP network from the device to a specific destination host.

```
7100 Chassis(rw)->tracert 192.167.252.17
tracert to 192.167.252.17 (192.167.252.17), 30 hops max, 40 byte packets
 1 matrix.extremenetworks.com (192.167.201.40)  20.000 ms  20.000 ms  20.000 ms
 2  14.1.0.45 (14.1.0.45)  40.000 ms  10.000 ms  20.000 ms
 3  192.167.252.17 (192.167.252.17)  50.000 ms  0.000 ms  20.000 ms
```

Use the **nslookup** command, in any command mode, to query name servers, translating hostnames to IP addresses or IP addresses to hostnames.

```
7100 Chassis(su)->nslookup -x 127.0.0.1
Name: localhost
Address: 127.0.0.1
```

## Switch Connection Statistics

Switch connection statistics can be displayed for:

- ICMP
- IP
- TCP
- UDP

Use the **show netstat** command to display switch connection statistics. Use the **stats** option to display statistics for all supported protocols.

```
7100 Chassis(rw)->show netstat stats
Ip:
 26034 total packets received
 25824 with invalid addresses
 0 forwarded
 0 incoming packets discarded
 187 incoming packets delivered
 6391 requests sent out
 21 dropped because of missing route
Icmp:
 14 ICMP messages received
 0 input ICMP message failed
ICMP input histogram:
   destination unreachable: 14
 6184 ICMP messages sent
 0 ICMP messages failed
ICMP output histogram:
   destination unreachable: 1
```



```

        echo request: 6183
    Tcp:
        2 active connection openings
        2 passive connection openings
        0 failed connection attempts
        0 connection resets received
        4 connections established
        153 segments received
        153 segments send out
        0 segments retransmitted
        0 bad segments received
        0 resets sent
    Udp:
        42 packets received
        1 packets to unknown port received
        0 packet receive errors
        57 packets sent
    7100 Chassis(rw)->
    
```

## Users

The network monitoring feature supports the display of information about the active console port or Telnet session(s) logged in to the switch. It also provides for the ability to send a message to one or all users with active sessions.

Use the **show users** command to display information for active console port or Telnet sessions on the switch.

```

    7100 Chassis(rw)->show users
        Session  User  Location
        -----  ----  -
    * console  admin console (via com.1.1)
      telnet   rw    134.141.192.18
    7100 Chassis(rw)->
    
```

Use the **tell** command to send a message to one or all users on the switch.

```

    7100 Chassis(rw)->tell rw@134.141.192.18 "System reset in 15 minutes"
    User rw@134.141.192.18 will receive:
    Message from admin@console: "System reset in 15 minutes"
    
```

## RMON

RMON (Remote Network Monitoring) is an industry standard specification that provides comprehensive network fault diagnosis, planning, and performance tuning information and allows for interoperability between SNMP management stations and monitoring agents. RMON extends the SNMP MIB capability by defining additional MIBs that generate a much richer set of data about network usage. These MIB “groups” each gather specific sets of data to meet common network monitoring requirements.

RMON statistics and history can be configured to gather packet counts in both the receive and transmit, receive only, or transmit only directions.

Table 26-1 lists:

- The four RMON monitoring groups supported on 7100-Series devices
- Each group's function
- The elements it monitors
- The group's associated commands

**Table 26-1 RMON Monitoring Group Functions and Commands**

RMON Group	What It Does...	What It Monitors...	CLI Command(s)
Statistics	Records statistics measured by the RMON probe for each monitored interface on the device.	Packets dropped, packets sent, bytes sent (octets), broadcast and multicast packets, CRC errors, oversized and undersized packets, fragments, jabbers, and counters for packets.	<b>show rmon stats</b> <b>set rmon stats</b> <b>clear rmon stats</b>
History	Records periodic statistical samples from a network.	Sample period, number of samples and item(s) sampled.	<b>show rmon history</b> <b>set rmon history</b> <b>clear rmon history</b>
Alarm	Periodically gathers statistical samples from variables in the probe and compares them with previously configured thresholds. If the monitored variable crosses a threshold, an event is generated.	Alarm type, interval, starting threshold, stop threshold.	<b>show rmon alarm</b> <b>set rmon alarm properties</b> <b>set rmon alarm status</b> <b>clear rmon alarm</b>
Event	Controls the generation and notification of events from the device.	Event type, description, last time event was sent.	<b>show rmon event</b> <b>set rmon event properties</b> <b>set rmon event status</b> <b>clear rmon event</b>

## SMON Priority and VLAN Statistics Counting

SMON is a set of RMON MIB extensions for switch monitoring. The 7100-Series supports the enabling and display of SMON Ethernet priority and VLAN statistics counters. SMON is described by RFC 2613. An SMON session for a specified port or range of ports is first created and then enabled before statistics are collected.

The 7100-Series platform supports a maximum of 16 SMON switch-wide VLAN sessions (using VTAP ports) and 28 SMON priority sessions. Port-VLAN sessions (for physical ports and LAGs) are not supported. Resources available to SMON priority and VLAN sessions are shared with other SMON tasks and port mirroring. Packets that match multiple SMON sessions may be counted in only one of them. Depending upon your configuration needs, you may not be able to configure the maximum number of supported SMON VLAN or priority sessions.

Use the **set smon priority create** and **set smon vlan create** commands to create priority and VLAN SMON sessions for the specified port(s) on the switch.

Use the **set smon priority enable** and **set smon vlan enable** commands to enable existing priority and VLAN SMON sessions for the specified port(s) on the switch.

The following example:

- Creates an SMON priority session for port ge.1.1
- Enables SMON priority for port ge.1.1
- Displays statistics for priority 0 for the enabled port:

```
7100 Chassis(rw)->set smon priority create ge.1.1
7100 Chassis(rw)->set smon priority enable ge.1.1
7100 Chassis(rw)->show smon priority ge.1.1 priority 0
Show Priority Statistics
```

```
-----
Interface = ge.1.1
Owner      = none
Creation   = 0 days 0 hours 37 minutes 36 seconds
Status     = enabled
-----
```

Priority 0	Packets	Octets
-----		
Total	3477	256168
Overflow	0	0

```
7100 Chassis(rw)->
```

The following example:

- Creates an SMON VLAN session for VLAN vtap.0.1
- Enables SMON VLAN monitoring for VLAN vtap.0.1
- Displays statistics for VLAN 1:

```
7100 Chassis(rw)->set smon vlan create vtap.0.1
7100 Chassis(rw)->set smon vlan enable vtap.0.1
7100 Chassis(rw)->show smon vlan 1
```

```
Show VLAN Statistics
-----
Interface = vtap.0.1
Owner      = none
Creation   = 20 days 1 hours 44 minutes 27 seconds
Status     = enabled
-----
```

VLAN 1	Packets	Octets
Total	3728	433041
Overflow	0	0
NonUnicast	2660	174336
NonUnicast Overflow	0	0

7100 Chassis (rw) ->

## Configuring Network Monitoring

This section provides details for the configuration of network monitoring on the 7100-Series products.

Table 26-2 lists network monitoring parameters and their default values.

**Table 26-2 Default Network Monitoring Parameters**

Parameter	Description	Default Value
history buffer	The number of lines of CLI input that are placed in a buffer for redisplay.	20 lines
buckets	The number of RMON history entries to maintain.	50 entries
interval	The period between RMON history or alarm sampling.	history = 1800 seconds alarm = 3600 seconds
owner	The RMON management station entity for a statistics or alarm context.	monitor
type	The RMON alarm monitoring method or property, RMON event, or TopN counter type.	alarm = absolute event = none topN = inpackets
startup	The RMON alarm type generated when an event is first enabled.	rising
rthresh	The RMON minimum threshold for causing a rising alarm.	0 events
fthresh	The RMON maximum threshold for causing a falling alarm.	0 events
revent	The RMON index event number to be triggered when the rising threshold is crossed.	0
fevent	The RMON index event number to be triggered when the falling threshold is crossed.	0
alarm or event status	Whether an entry is enabled or disabled.	disabled

To optionally change the size of the history buffer, use the **set history** command, specifying the size of the history buffer. The **default** option configures the specified history buffer setting to persist for all future sessions. Otherwise, the setting only affects this session.

This example shows how to set the size of the command history buffer to 25 lines and make this the default setting:

```
7100 Chassis (rw) -> set history 25 default
```

To optionally send a message to one or all active users on this switch, use the **tell** command, specifying an individual destination or all users. The **dest** option specifies the user and location in the user@location format.

This example shows how to tell user rw@134.141.192.18 about a system reset:

```
7100 Chassis(rw)->show users
  Session  User  Location
  -----  ----  -----
* console  admin console (via com.1.1)
  telnet   rw    134.141.192.18
7100 Chassis(rw)->tell rw@134.141.192.18 "System reset in 15 minutes"
```

User rw@134.141.192.18 will receive:

```
Message from admin@console: "System reset in 15 minutes"
```

[Table 26-3](#) describes network diagnostics commands.

**Table 26-3 Network Diagnostics Commands**

Task	Command
<p>To determine the availability of another node on the network:</p> <ul style="list-style-type: none"> <li>• <b>-s bytes</b> – (Optional) Specifies the number of data bytes to be sent.</li> <li>• <b>-c count</b> – (Optional) Number of ping packets.</li> <li>• <b>-n</b> – (Optional) Avoids any communications with nameservers.</li> <li>• <b>-p pattern</b> – (Optional) Specify up to a 16 bit hexadecimal pattern to fill outgoing packet with (ex. -p ff).</li> <li>• <b>-t hops</b> – (Optional) Specifies the maximum number of hops for the ping.</li> <li>• <b>-I interface</b> – (Optional) Source IP Interface.</li> <li>• <b>-S ip-address</b> – (Optional) Source IP address.</li> <li>• <b>-Q service-type</b> – (Optional) Specifies the Type of Service in the IPv4 header or the traffic class in the IPv6 header.</li> <li>• <b>-r</b> – (Optional) Bypass the normal routing tables and send directly to a host on an attached network.</li> <li>• <b>-i</b> – (Optional) Specifies the time in milliseconds to wait for ping timeouts and between sending ping packets.</li> <li>• <b>-v</b> – (Optional) Forces ping to a specific ip version.</li> <li>• <b>-V router</b> – (Optional) Specify a virtual router name or number for this ping.</li> </ul> <p><i>host</i> – Specifies the IP address or a hostname of the receiving device.</p>	<p><b>ping</b> [-s bytes] [-c count] [-n] [-p pattern] [-t milliseconds] [-I interface] [-S ip-address] [-Q service-type] [-r] [-i milliseconds] [-v {4   6}] [-V router] host</p>
<p>To display a hop-by-hop path from the device to a specific destination host:</p> <ul style="list-style-type: none"> <li>• <b>-d ip-address</b> – (Optional) Performs a reverse lookup (finds a hostname that matches the specified IP address).</li> <li>• <b>-F</b> – (Optional) Specifies that the traceroute packet should not be fragmented.</li> <li>• <b>-f first-TTL</b> – (Optional) Specifies the maximum Time-To-Live (TTL) used in the first outgoing probe packets.</li> </ul>	<p><b>traceroute</b> [-d ip-address] [-F] [-f first_ttl] [-I] [-i interface] [-m max_ttl] [-n] [-p port] [-q nqueries] [-r] [-s source-address] [-t tos] [-v {4   6}] [-V router] [-w waittime] [-x] host</p>

**Table 26-3 Network Diagnostics Commands (continued)**

Task	Command
<p><b>-I</b> – (Optional) Specifies that ICMP echo requests should be used instead of UDP datagrams.</p> <p><b>-i source-interface</b> – (Optional) Specifies the IP source interface (for example <code>vlan.0.5</code> for VLAN 5).</p> <p><b>-m max-ttl</b> – (Optional) Specifies the maximum Time-To-Live (TTL) for outgoing packets.</p> <p><b>-n host-ip-address</b> – (Optional) Specifies that name server contact should be avoided.</p> <p><b>-p udp-dest-port</b> – (Optional) Specifies the initial UDP destination port. For each sent probe the UDP destination port is increased by one.</p> <p><b>-q number-of-probes</b> – (Optional) Specifies the number of probes to send out for each hop.</p> <p><b>-r</b> – (Optional) Specifies that normal host routing tables should be bypassed.</p> <p><b>-s source-ip-address</b> – (Optional) Specifies the source IP address for the traceroute probes.</p> <p><b>-t tos</b> – (Optional) Specifies the Type-of-Service (ToS) for IPv4 or the traffic class for IPv6.</p> <p><b>-v version</b> – (Optional) Forces traceroute to use either IPv4 or IPv6.</p> <p><b>-V router</b> – (Optional) Specifies the virtual router to use for this traceroute.</p> <p><b>-w period</b> (Optional) Specifies the time in seconds to wait for a response to a probe.</p> <p><b>-x</b> – (Optional) Specifies that traceroute should not calculate checksum.</p> <p><b>host host</b> – Specifies an IP address or a host to find a route to.</p>	
<p>To query a name server to translate hostnames to IP addresses or IP addresses to hostnames:</p> <p><b>-x</b> – (Optional) Specifies that a reverse lookup should be performed. If this parameter is used, then you must specify an IP address as the host variable.</p> <p><b>-v {4   6}</b> – (Optional) Specifies the IP version for this name server lookup.</p> <p><b>host</b> – Specifies the host name, or an IP address, in the case of a reverse lookup.</p>	<b>nslookup [-x] [-v {4   6}] host</b>

[Procedure 26-1](#) describes how to configure SMON. SMON commands can be entered from any command mode.

### Procedure 26-1 Configuring SMON

Step	Task	Command(s)
1.	Optionally, first create and then enable an SMON session for the collection of Ethernet priority statistics for the specified port(s).	<b>set smon priority {create   enable} port-string [owner]</b>

**Procedure 26-1 Configuring SMON (continued)**

Step	Task	Command(s)
2.	Optionally, first create and then enable an SMON session for the collection of VLAN statistics for the selected VLANs, specified using VTAP port strings. (Port-VLAN collections for individual Ethernet and LAG ports are not supported).	<b>set smon vlan</b> {create   enable} <i>port-string</i> [owner]

[Procedure 26-2](#) describes how to configure RMON. RMON commands can be entered from any command mode.

**Procedure 26-2 Configuring Remote Network Monitoring**

Step	Task	Command(s)
1.	Optionally, configure RMON to create entries that record statistics measured by the RMON probe for each specified interface. <ul style="list-style-type: none"> <li><i>index</i> – Specifies the index number for this entry</li> <li><i>port-string</i> – assigns this entry to a specific port</li> <li><b>owner</b> – (Optional) Specifies the management station owner for this entry</li> <li><b>direction</b> – (Optional) Specifies the direction in which RMON statistics are collected</li> </ul>	<b>set rmon stats</b> <i>index</i> [ <i>port-string</i> ] [owner] [direction {rx+tx   rx   tx}]
2.	Optionally, specify the maximum number and period for recorded statistical samples from a network. <ul style="list-style-type: none"> <li><i>index</i> – Specifies the index number for this entry</li> <li><i>port-string</i> – assigns this entry to a specific port</li> <li><b>bucket</b> – (Optional) Specifies the maximum number of entries to maintain</li> <li><b>interval</b> – (Optional) Specifies the period between samples in seconds</li> <li><b>owner</b> – (Optional) Specifies the management station owner for this entry</li> <li><b>direction</b> – (Optional) specifies the direction in which RMON history is collected</li> </ul>	<b>set rmon history</b> <i>index</i> [ <i>port-string</i> ] [buckets buckets] [interval <i>interval</i> ] [owner <i>owner</i> ] [direction {rx+tx   rx   tx}]
3.	Configure RMON probe variable thresholds that will trigger an alarm if crossed by a sampled probe. <ul style="list-style-type: none"> <li><i>index</i> - Specifies the entry for this set of alarm properties</li> <li><b>interval</b> - (Optional) Specifies the period between samples in seconds</li> <li><b>object</b> - (Optional) Specifies the MIB object to be monitored</li> <li><b>type</b> - (Optional) Specifies a monitoring method</li> </ul>	<b>set rmon alarm properties</b> <i>index</i> [interval <i>interval</i> ] [object <i>object</i> ] [type {absolute   delta}] [startup {rising   falling   either}] [rthresh <i>rthresh</i> ] [fthresh <i>fthresh</i> ] [revent <i>revent</i> ] [fevent <i>fevent</i> ] [owner <i>owner</i> ]

**Procedure 26-2 Configuring Remote Network Monitoring (continued)**

Step	Task	Command(s)
	<ul style="list-style-type: none"> <li>• <b>startup</b> - (Optional) Specifies the alarm type generated when this event is first enabled</li> <li>• <b>rthresh</b> - (Optional) Specifies the minimum threshold that will cause a rising alarm</li> <li>• <b>fthresh</b> - (Optional) Specifies the minimum threshold that will cause a falling alarm</li> <li>• <b>revent</b> - (Optional) Specifies the index number of the RMON event to be triggered when the rising threshold is crossed</li> <li>• <b>fevent</b> - (Optional) Specifies the index number of the RMON event to be triggered when the falling threshold is crossed</li> <li>• <b>owner</b> - (Optional) Specifies the management station owner for this entry</li> </ul>	
4.	Enable a configured alarm entry.	<b>set rmon alarm status</b> <i>index</i> <b>enable</b>
5.	Configure RMON probe variable thresholds that will trigger an event if crossed by a sampled probe. <ul style="list-style-type: none"> <li>• <i>index</i> - Specifies the entry for this set of event properties</li> <li>• <b>description</b> - (Optional) Specifies a text string description for this event</li> <li>• <b>type</b> - (Optional) Specifies the event notification type for this entry</li> <li>• <b>community</b> - (Optional) Specifies an SNMP community name to use if the message type is set to trap</li> <li>• <b>owner</b> - (Optional) Specifies the management station owner for this entry</li> </ul>	<b>set rmon event properties</b> <i>index</i> [ <b>description</b> <i>description</i> ] [ <b>type</b> { <i>none</i>   <i>log</i>   <i>trap</i>   <i>both</i> }] [ <b>community</b> <i>community</i> ] [ <b>owner</b> <i>owner</i> ]
6.	Enable a configured event entry.	<b>set rmon event status</b> <i>index</i> <b>enable</b>

Table 26-4 describes how to manage network monitoring.

**Table 26-4 Managing Network Monitoring**

Task	Command
To disconnect from a console or Telnet session:	<b>disconnect</b> { <i>ip-address</i>   <b>console</b> }
To disable SMON priority counters collection for the specified port(s), without clearing the created session:	<b>set smon priority disable</b> <i>port-string</i> [ <i>owner</i> ]
To disable SMON VLAN counters collection for the specified selected VLANs, specified using VTAP port strings, without clearing the created session:	<b>set smon vlan disable</b> <i>port-string</i> [ <i>owner</i> ]
To clear an existing SMON priority counters session for the specified port(s):	<b>clear smon priority</b> [ <i>port-string</i> ]
To clear an existing SMON VLAN counters session for the specified port(s):	<b>clear smon vlan</b> [ <i>port-string</i> ]
To delete one or more RMON statistics entries:	<b>clear rmon stats</b> { <i>index</i>   <b>to-defaults</b> }



**Table 26-4 Managing Network Monitoring (continued)**

Task	Command
To delete one or more RMON statistics entries:	<b>clear rmon stats</b> <i>{index-list   to-defaults}</i>
To delete one or more RMON history entries:	<b>clear rmon history</b> <i>{index-list   to-defaults}</i>
To delete an RMON alarm entry:	<b>clear rmon alarm</b> <i>index</i>
To delete an RMON event entry and any associated log entries:	<b>clear rmon event</b> <i>index</i>

[Table 26-5](#) describes how to display network monitoring information and statistics.

**Table 26-5 Displaying Network Monitoring Information and Statistics**

Task	Command
To display the contents of the CLI history buffer:	<b>history</b>
To display the current history buffer size setting:	<b>show history</b>
To display switch connection statistics for all or the specified protocol:	<b>show netstat</b> [ <i>icmp   ip   stats   tcp   udp</i> ]
To display information for the active console port or Telnet sessions on the switch:	<b>show user</b>
To display SMON priority statistics counters for all or the specified port(s) and priorities:	<b>show smon priority</b> [ <i>port-string</i> ] [ <i>priority priority</i> ]
To display SMON VLAN statistics counters for all or the specified VLAN(s):	<b>show smon vlan</b> [ <i>port-string</i> ] [ <i>vlan vlan-id</i> ]
To display RMON statistics for one or more ports:	<b>show rmon stats</b> [ <i>port-string</i> ] [ <i>wide</i> ] [ <i>by-size</i> ]
To display RMON history properties and statistics:	<b>show rmon history</b> [ <i>port-string</i> ] [ <i>wide</i> ] [ <i>interval {30sec   5min   25min}</i> ]
To display RMON alarm entries:	<b>show rmon alarm</b> [ <i>index</i> ]
To display RMON event entry properties:	<b>show rmon event</b> [ <i>index</i> ]

## Connectivity Fault Management Configuration

This chapter provides information about configuring and monitoring the Connectivity Fault Management (CFM) protocol on 7100-Series devices.

For information about...	Refer to page...
<a href="#">How to Use Connectivity Fault Management in Your Network</a>	<a href="#">27-1</a>
<a href="#">Connectivity Fault Management Overview</a>	<a href="#">27-2</a>
<a href="#">Implementing Connectivity Fault Management</a>	<a href="#">27-8</a>
<a href="#">Configuring CFM at the Global System Level</a>	<a href="#">27-9</a>
<a href="#">Activating CFM Configuration</a>	<a href="#">27-11</a>
<a href="#">Configuring a Maintenance Domain (MD)</a>	<a href="#">27-11</a>
<a href="#">Configuring a Maintenance Association (MA)</a>	<a href="#">27-15</a>
<a href="#">Configuring a Maintenance End-Point (MEP)</a>	<a href="#">27-18</a>
<a href="#">Configuring Connectivity Fault Management</a>	<a href="#">27-24</a>
<a href="#">Terms and Definitions</a>	<a href="#">27-44</a>

### How to Use Connectivity Fault Management in Your Network

The Connectivity Fault Management (CFM) process, as defined in the IEEE 802.1Q-2011 standard, provides network operators the means to monitor and troubleshoot services that may span multiple domain Ethernet networks. It provides a set of diagnostics and monitoring functions at the service provider level, allowing operators:

- To determine the relative health of an end-to-end network service and operational status
- To identify faults or mis-configurations within a network
- To take administrative action to correct those issues

CFM allows a customer to validate the end-to-end Ethernet service within a single domain network, a large network that is segmented into separate domains within the same organization, or a network that has contracted with a service provider. In the event of a degradation of that service, CFM provides diagnostic data that can locate the problem within a single organization or that can be forwarded to the service provider. Whether working with an outside entity or within a single organization, CFM can:

- Determine that the end-to-end service is properly configured
- Validate that all network nodes intended to be attached to the end-to-end service are reachable using that service

- Help pinpoint where a loss in connectivity may be located

The Extreme Networks 7100-Series CFM implementation supports the monitoring of a VLAN service, and the association of one or more VLANs with the primary service. For the remaining discussion in this chapter, any reference to a CFM monitored service is a reference to a CFM monitored VLAN service. A CFM service is monitored by the periodic sending of continuity check messages (CCM) across the monitored service. A CCM is a multicast message, confined to a single operator domain that provides a means to detect connectivity failures or configuration errors for that monitored service. These messages are unidirectional and do not solicit a response. Each end of a CFM monitored service transmits a periodic multicast CCM inward towards the other end of the service. The primary service can both send and receive CCMs. Services associated with the primary can only receive CCMs.

A CFM monitored service resides within three hierarchical layers of CFM configuration:

- Maintenance Domain (MD) – A logical container for all the equipment associated with the CFM monitored service and owned by a single network operator.
- Maintenance Association (MA) – A logical container for a specific CFM monitored service.
- Maintenance Point (MP) – A demarcation point on a port that implements the CFM functions within an MA. There are two types of MPs: Maintenance End-Points (MEP) and Maintenance Intermediate-Points (MIP)
  - MEPs – Ports, belonging to an MA, through which data enters and exits the portion of the network monitored by the CFM service.
  - MIPs – Auto-created MPs on ports that reside along the path between MEPs. The MIP supplements the function to the MEPs of the domain by passively snooping the CCMs that pass through them.

Multiple services can be associated with the primary service by configuring a VLAN table for the primary service.

The CFM loopback protocol provides for connectivity verification by sending loopback messages between the initiating device and a MEP or a MAC address. The CFM linktrace protocol provides for path verification and helps identify where in the path a connectivity problem is located. Linktrace messages can be sent between the initiating device and a MEP or a MAC address. CFM can notify the network operator by Syslog or SNMP traps when connectivity failures or configuration errors are detected.

The minimum generated alarm defect can be set for both Syslogs and traps. See “[MEP Defect Definitions](#)” on page 27-21 for reported defect details. Logging can be filtered by MD, MA, or MEP.

CFM configuration takes place within an hierarchy of configuration modes or contexts (see “[CFM Configuration Modes](#)” on page 27-8):

- The system global configuration mode
- The global system and CFM service specific default modes
- The three configuration layers in which a CFM service resides: MD, MA, and MEP
- The MA component configuration context which is a subset of the MA configuration

## Connectivity Fault Management Overview

For information about...	Refer to page...
<a href="#">Maintenance Domain (MD)</a>	<a href="#">27-3</a>

For information about...	Refer to page...
<a href="#">Maintenance Association (MA)</a>	27-4
<a href="#">Maintenance Point (MP)</a>	27-5
<a href="#">CFM Configuration Modes</a>	27-8

## Maintenance Domain (MD)

A maintenance domain (MD) is the highest configuration context in which a CFM service resides. The MD is a collection of network devices, typically owned and operated by a single organization. Management of devices within a domain falls under the control of that single organization. Domains must be contiguous, such that all the devices belonging to a domain have uninterrupted network connectivity with each other. Domains may be nested or adjacent, but can not share network devices with other domains. Domains are intended to provide connectivity to systems outside of the domain. All devices associated with a monitored service must belong to the MD. The monitored service is configured within the MD. MDs are uniquely identified by an MD Name.

An MD is assigned a level relative to other MDs a monitored service passes through. The monitored service must belong to an MD with a higher level than all other MDs that the monitored service passes through. The CCM will not be allowed to enter an MD with a higher level than the MD that initiated the CCM.

Figure 27-1 displays how each administratively controlled segment of the network is configured within its own MD in a service provider context made up of:

- A customer equipment maintenance domain **mdCE1**
- A service provider maintenance domain **mdSP1**
- Two operator maintenance domains **mdOp1** and **mdOp2**

**Figure 27-1 Maintenance Domain Overview**

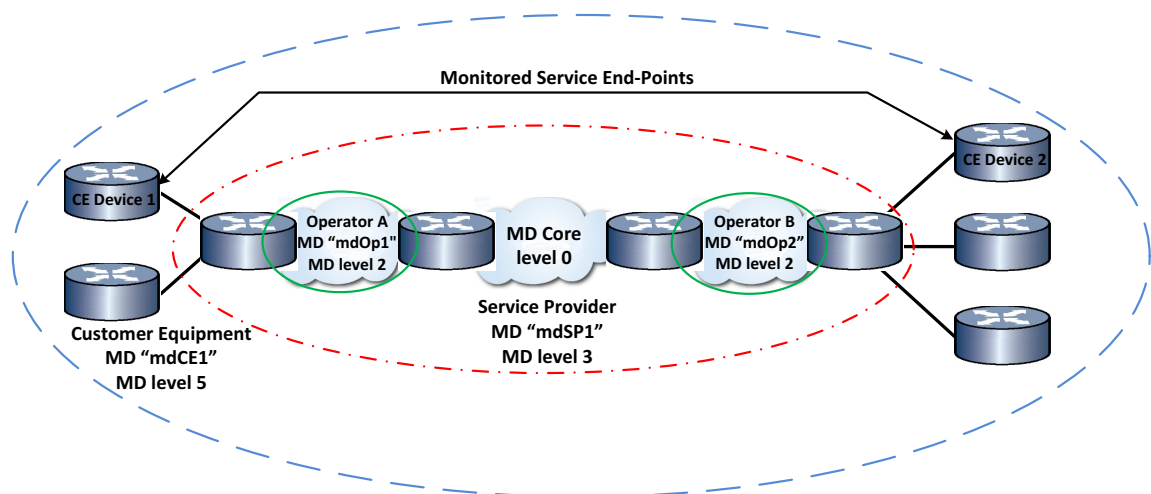


Figure 27-1 displays a typical MD domain configuration in a service provider context. The end-points of the monitored service (CE Device 1 and CE Device 2) belong to the customer equipment domain **mdCE1**. The monitored service is a VLAN these end-points belong to. The MD to which the monitored service devices belong must encapsulate any MDs the monitored service passes through from one end-point of the service to the other end-point. This encapsulation is

accomplished by the MD level. The customer equipment MD is assigned the highest MD level in our example, assuring that CCMs will be allowed to pass between the monitored service end-points. The encapsulated MDs can belong to such entities as service providers and network operators as shown here, or they can be segmented parts of your own network.

In our example, the monitored service belongs to MD **mdCE1** and passes through **mdSP1**, **mdOp1**, and **mdOp2** before reaching end-point CE Device 2 which belongs to MD **mdCE1**.

An MD constrains CFM traffic flows. CFM traffic flows of the MD owner and can transparently flow through any encapsulated MD with a lower MD level. mdCE1 CFM traffic flows pass through all the displayed MDs. mdSP1 traffic flows pass through mdOp1 and mdOp2, but are prevented from transiting into mdCE1. mdOP1 and mdOP2 CFM traffic is restricted to its own MD. This containment of traffic flows prevents unintended information flow between MDs.


## Maintenance Association (MA)

A maintenance association (MA) uniquely identifies a service within an MD. A service may be defined by an individual primary VLAN, and optionally by one or more VLANs associated with the primary VLAN, using a VLAN table configuration. There may be multiple MAs within a domain. Subsets of devices residing within the domain are collectively configured to form these associations. The devices belonging to a particular association will communicate among each other to implement the various features provided by CFM.

[Figure 27-2](#) displays three maintenance associations configured for the customer equipment level 5 MD **mdCE1**. In this presentation, MD levels 0, 2, and 3 as shown in [Figure 27-1](#) on page 27-3 are abstracted in the cloud. The presence of the non-customer equipment MDs in the cloud is transparent to the customer equipment MAs. Each network device on the edge of the cloud represents a CFM networking device configured for the appropriate MA that monitors the primary VLAN the device is on:

- VLAN 100 – MA: **maCE1**
- VLAN 200 – MA: **maCE2**
- VLAN 300 – MA: **maCE3**

Figure 27-2 Maintenance Association Overview



For each monitored VLAN, each network device depicted resides at the edge of the MD for its physical location; for example it may be on a separate university or enterprise campus from the other nodes of the monitored service. Data from these nodes may or may not have to transit an ISP or other operator domains outside of the customer equipment administrator's control, before reaching the other end of the monitored VLAN.

Within MA configuration you specify the monitored service (VLAN), create a list of all MEPs that belong to the MA and optionally change the interval between the sending of CCMs.

## Maintenance Point (MP)

A maintenance point (MP) is a demarcation point on a port that implements the CFM functions within an MA. MPs serve as points of contact allowing communication with other devices within the MA. MPs also operate as filters to confine CFM Ethernet frames within the bounds of a domain by dropping frames that do not belong to the correct level. There are two types of MPs: maintenance end-points (MEPs) and maintenance intermediate points (MIPs).

### Maintenance End-Point (MEP)

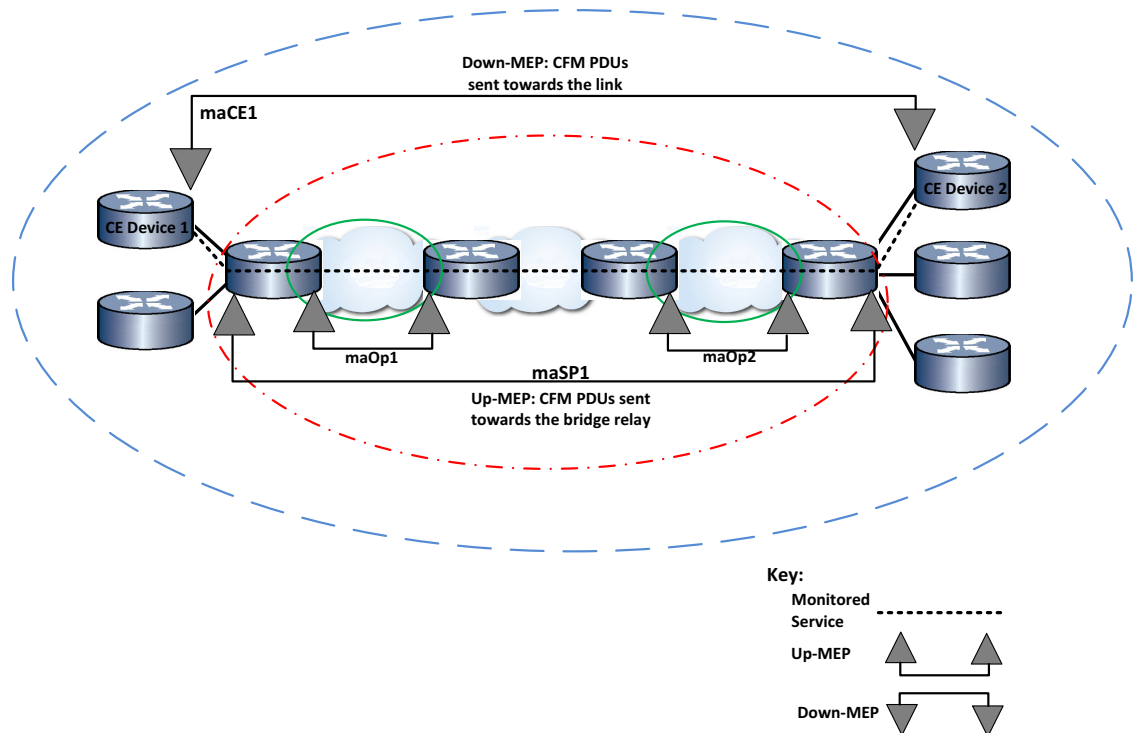
The Maintenance End-Point (MEP) serves as the logical boundary between devices that belong to different maintenance domains. A MEP resides at the edge of an MD. A MEP is associated with a single MA that monitors a primary VLAN service and any services associated with the primary service configured in an enabled VLAN table. The MEP must belong to the primary service associated with the MA or a service associated with the primary service. A MEP has a direction of either Up or Down. A Down-MEP sends CFM PDUs towards and receives CFM PDUs from the

link. In our overview example, The MEPs associated with the customer equipment monitored service configure their MEPs as Down-MEPs. An Up-MEP sends PDUs towards the bridge relay and receives PDUs from the bridge relay. Up-MEPs communicate through the bridge relay with all other ports attached to the protective service within that bridge. All non-customer equipment MEPs are configured as Up-MEPs for a customer equipment monitored service.

Figure 27-3 on page 27-6 displays MEP location and direction for our overview example. Customer equipment devices CE Device 1 and CE Device 2 belong to MA **maCE1** and each are configured with a Down-MEP facing the service provider MD configured as MA **maSP1**. The customer equipment administrator wants to monitor the service between CE Device 1 and CE Device 2. The customer equipment Down-MEPs send CFM PDUs towards the link and therefore towards the remote Down-MEP belonging to **maCE1** at the other end of the service. All Up-MEPs, belonging to other domains configured between the initiating and remote customer equipment down-MEPs, will transparently forward the CFM PDUs to the remote customer equipment MEP.

If a MEP fails to receive CFM PDUs in a timely fashion, CFM reports the failure.

**Figure 27-3 Maintenance End-Point Overview**



Within MEP configuration mode you must specify the port that the MEP is on and its direction. The MEP direction defaults to down.

## Maintenance Intermediate Point (MIP)

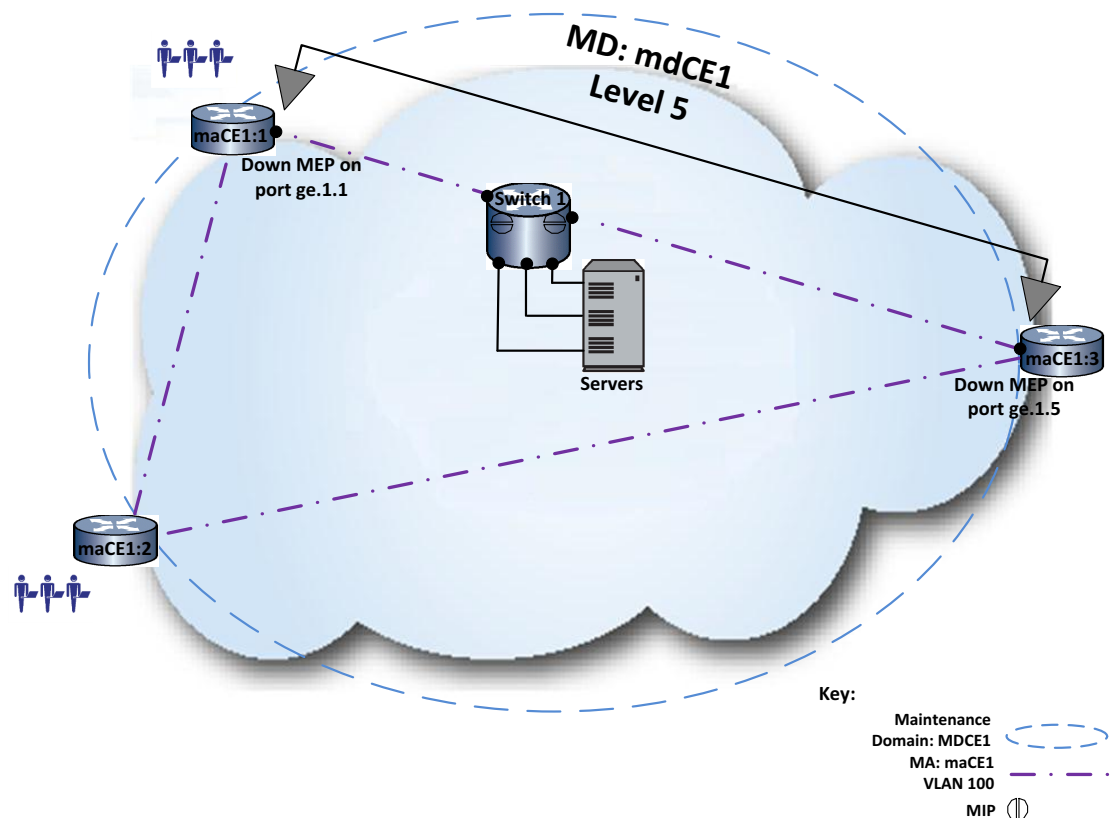
A Maintenance Intermediate Point (MIP) resides in the interior of an MD. MIPs are created on ports that reside along the path between MEPs. The MIP supplements the function to the MEPs of the domain. MIPs passively collect information by snooping the CCMs that pass through them. The information is collected in a database. These MIP databases act as highway “mile-markers” along the continuity check message path. MIPs may respond to loopback and linktrace requests received from MEPs in its MD. Use the loopback protocol to determine that a problem exists. Use the linktrace protocol to determine which MIP can identify the problem.

MIPs are comprised of two Maintenance Half Functions (MHFs), one with an “up” direction, and one with a “down” direction. The up direction half function points towards the bridge relay function. The down direction half function points towards the link. Use the **show cfm stack-table** command to access MIP direction port information. MIPs for a given service can be configured in any domain but must be configured with the same MD level as the MD of the MEPs sending the continuity check messages. MIPs can not be created on devices containing down-MEPs. MIPs can be created on devices containing Up-MEPs or no MEPs. MIPs can be created on down-MEPs using MD default configuration.

You do not administratively configure MIPs. MIPs are automatically created if MHF creation is enabled. If MHF creation is not enabled, MIPs are not created.

Figure 27-4 on page 27-7 presents a typical situation in which you would want to turn on MIPs. Switch 1 resides in the interior of the MD on a VLAN monitored by MA maCE1. Down-MEPs are configured on port ge.1.1 of maCE1:1 and port ge.1.5 of device maCE1:3. The MD default setting for MHF creation has been enabled on Switch 1. The ports on Switch 1 used by the service are being passively monitored by MIPs. If a problem is reported on one of the servers connected to Switch 1, from a CFM enabled device, use the loopback protocol specifying the MAC address of the Switch 1 port to verify connectivity with the server connected port. Use the linkstate protocol specifying the MAC address of the Switch 1 port to verify that a path exists between the initiating device and the Switch port. Because the server connected ports are on the edge of the network, you may want to put down MEPs on those ports.

**Figure 27-4 Maintenance Intermediate-Point Overview**





## CFM Configuration Modes

[Table 27-1](#) lists the seven modes in which CFM configuration takes place.

**Table 27-1 CFM Configuration Modes**

Configuration Mode	Description
Global System configuration	System level CFM configuration and access to CFM default and MD configuration modes takes place in global system configuration mode. See <a href="#">Table 27-5</a> on page 27-25 for a listing of CFM global system configuration commands.
Default Maintenance Domain configuration	Default maintenance domain is the CFM configuration mode with the lowest precedence. If no other configuration levels are administratively configured for a given global CFM parameter, the system default maintenance domain configured value for that parameter is used. Throughout this document system default maintenance domain configuration mode is referred to as system default MD. See <a href="#">Table 27-6</a> on page 27-26 for a listing of system default MD configuration commands.
Default maintenance domain VLAN	Default maintenance domain VLAN configuration mode has a higher precedence than system default MD configuration but is lower than all other CFM modes that contain the same parameters. See <a href="#">Table 27-6</a> on page 27-26 for a listing of Default maintenance domain VLAN configuration commands.
Maintenance Domain (MD) configuration	Provides access to the MD configuration commands and MA configuration mode. Through out this document maintenance domain configuration mode is referred to as MD configuration mode. MD configuration has a higher precedence than either of the default configuration modes. See <a href="#">Procedure 27-1</a> on page 27-26 for a listing of maintenance domain configuration commands.
Maintenance Association (MA) configuration	Provides access to the MA configuration commands and MA-Comp and MEP configuration modes. Through out this document maintenance association configuration mode is referred to as MA configuration mode. See <a href="#">Procedure 27-2</a> on page 27-27 for a listing of maintenance association configuration commands.
Maintenance Association Component (MA-Comp) configuration	Provides access to MA component configuration commands. MA component configuration is a subset of MA configuration. Through out this document maintenance association component configuration mode is referred to as MA-Comp configuration mode. MA-Comp configuration has a higher precedence than either of the default modes and the MD configuration mode. See <a href="#">Procedure 27-3</a> on page 27-27 for a listing of MA component configuration commands.
Maintenance End-Point (MEP) configuration	Provides access to MEP configuration commands. Through out this document maintenance end-point configuration mode is referred to as MEP configuration mode. See <a href="#">Procedure 27-4</a> on page 27-28 for a listing of maintenance end-point configuration commands.

## Implementing Connectivity Fault Management

To Implement CFM on an Extreme Networks 7100-Series platform:

- Globally enable CFM on the device.
- Optionally modify system MD and MD VLAN defaults
- Optionally set the logging filter for Syslog messages

- Optionally configure VLAN tables that associate one or more CFM services with a primary service
- Configure the MDs for your system
- Configure an MA for each CFM monitored service
- Configure the MEP-list for each MA
- Configure the MEPs associated with each MA

## Configuring CFM at the Global System Level

Within the global configuration command mode, you:

- Globally enable CFM on the device
- Access the default modes for configuring global MD default and MD default VLAN service parameters (see “[CFM Configuration Modes](#)” on page 27-8)
- Access the MD configuration mode (see “[CFM Configuration Modes](#)” on page 27-8)
- Configure Syslog message filtering

CFM must be enabled globally for CFM to be operational. CFM is disabled by default. Use the **cfm enable** command in global configuration mode to globally enable CFM on the device.

### CFM Logging Filtering

By default, all CFM MD, MA, and MEP Syslog messages are sent. CFM logging filtering allows you to limit the sending of logged messages by a specified MD, MA or MEP. If an MA or MEP is not specified, Syslog messages are sent for all MAs and MEPs for the specified MD.

Use the **cfm logging filter** command to filter the sending of CFM Syslog messages to a specified MD, MA and MEP.

This example shows how to configure Syslog to display Syslog messages for all MEPs in the **myMA1** maintenance association of the **myMD1** maintenance domain:

```
7100 Chassis(rw-config)->cfm logging filter md string-name myMD1 ma string-name
myMA1
7100 Chassis(rw-config)->
```

### VLAN Table Configuration

By default there is a one-to-one relationship between a service ID and the CFM monitored service. If the service ID is 10 and the monitored CFM service is a VLAN, the monitored service is VLAN 10. In a one-to-one CFM service, the only monitored service is the primary service or selector. There are no other services associated with the primary service. The primary service can both receive and send CFM PDUs when the CFM device is a MEP. The primary service can receive CFM PDUs when the CFM device is a MIP. To fully monitor multiple services on the MEP would require that you create a separate service (MA) for each of the VLANs associated with a MEP adding considerable administrative overhead. When the primary goal is to fully monitor a single CFM service and, at the same time, monitor the reception of CFM PDUs on a group of associated services within a single MA, the VLAN table provides a significant reduction in administrative overhead.

The CFM VLAN table allows you to associate one or more services or selectors with the primary service. These services associated with the primary are capable of receiving CFM PDUs on any service configured in the VLAN table to which the monitored service belongs. For example: if the

VLAN table for primary selector 10 has a configured selector list containing selectors 11 - 14, and the monitored service is a VLAN, MEPs with any primary VLAN of 10 through 14 will:

- Both receive and send CFM PDUs on the MEP's primary VLAN (the configured VID or the primary service configured for the MA if the MEP VID is set to 0)
- Receive CFM PDUs on any other VLAN table member configured on the device


Figure 27-5 presents a VLAN table overview for an MA configured with a VLAN service and a VID of 10. The VLAN table is configured with a primary selector of 10 and associated selector list of 11 and 12. MEP0 and MEP2 are configured with the default VID 0 and inherit the MA VID setting 10. For these two MEPs, CFM PDUs are sent and received on VLAN 10 and received on VLANs 11 and 12. MEP1 and MEP3 are configured for VID 11. For these two MEPs, CFM PDUs are sent and received on VLAN 11 and received on VLANs 10 and 12. MEP4 is configured with VID 12. For MEP4, CFM PDUs are sent and received on VLAN 12 and received on VLANs 10 and 11.

Figure 27-5 VLAN Table Configuration Overview

**VLAN Table**

Primary Selector: 10

Selector List: 11, 12



The primary selector defines the ID of the service being modified. When configuring a CFM service, the configured VID for the new service is the primary selector. If the VID is configured for the default of 0, the MA the device belongs to provides the primary selector. If the device primary selector is a member of the VLAN table, all other members of the table are associated with the device primary selector.

When configuring the VLAN table, the primary selector is specified. One or more service IDs are specified in the selector list. Specifying the **enable** command option activates the association between the IDs in the selector list and the primary service. Specifying the **disable** command

option disables the association of the IDs in the selector list with the primary service. The VLAN table remains configured, but is not active.

Maintenance points (MIP or MEP) associated with a CFM service will be able to receive CFM PDUs on any of the active IDs defined in the VLAN table selector list. MEPs may be configured with a Primary VID, which must belong to the enabled VLAN table, otherwise there is a one-to-one relationship between the VID and the service. With a primary VID defined, the MEP can transmit tagged PDUs using that primary service. If no primary VID is defined by the maintenance point, the VID, as inherited from the MA, will be used by the maintenance point to tag its transmitted PDUs.

Use the `cfm vlan-table` command in global configuration mode to configure a VLAN table.

## Activating CFM Configuration

Changes made in the named MD, MA, MA component, and MEP configuration contexts do not take affect until the the CFM configuration is activated for that context. Use the **enable** command in the appropriate configuration context to activate the CFM configuration for that context.

If **enable** has already been entered in the current context, you must first enter **no enable** before making any further changes. If you attempt to make changes in a context that has already been enabled, you receive an error message like the following MD context error message:

```
Error: MD must be disabled ("no enable") before changes can be made.
```

Once changes are completed, enter **enable** again for the changes to take affect.

For a given configuration context to be operational, its parent context must be enabled:

- The MD context is the parent of MA and MA component contexts
- The MA context is the parent of the MA component context
- The MA and MA component contexts are the parents of the MEP context.

## Configuring a Maintenance Domain (MD)

Refer to “[Maintenance Domain \(MD\)](#)” on page 27-3 for an MD overview discussion.

For information about...	Refer to page...
<a href="#">MD Configuration Modes</a>	<a href="#">27-11</a>
<a href="#">MD Naming Conventions</a>	<a href="#">27-13</a>
<a href="#">Setting SenderID TLV Permission</a>	<a href="#">27-13</a>
<a href="#">Enabling Maintenance Intermediate-Points (MIP)</a>	<a href="#">27-14</a>
<a href="#">Setting the MD Level</a>	<a href="#">27-14</a>
<a href="#">Changing the Maintenance Domain Name</a>	<a href="#">27-15</a>

## MD Configuration Modes

You access MD configuration from global configuration mode. There are two CFM MD default configuration modes and a named MD configuration mode.

**Table 27-2 MD Configuration Modes**

Mode	Description
<b>Global MD and VLAN MD Service Default Configuration Mode</b>	<p>Global MD default configuration values are used when global default parameters have not been administratively modified in higher precedence configuration modes. Global MD default configuration mode is the lowest precedence configuration mode.</p> <p>VLAN MD service default configuration values are applied to the specified VLAN service. CFM monitors a network service. CFM currently supports VLAN monitoring. Any MD default parameters configured in MD VLAN service default mode take precedence over global MD default mode configuration for the specified VLAN service.</p> <p>In global MD and VLAN service default modes you can configure:</p> <p>ID permission – See <a href="#">“Setting SenderID TLV Permission”</a> on page 27-13</p> <p>MIP creation – See <a href="#">“Enabling Maintenance Intermediate-Points (MIP)”</a> on page 27-14</p> <p>MD level – See <a href="#">“Setting the MD Level”</a> on page 27-14</p> <p>Use the <b>cfm default-md default</b> command to enter global MD default configuration mode.</p> <p>Use the <b>cfm default-md vid</b> command, specifying the VLAN service to be monitored, to enter MD VLAN service default configuration mode.</p>
<b>Named MD Configuration Mode</b>	<p>Configuration entered in a named MD configuration mode takes precedence over global MD default configuration and MD VLAN service defaults.</p> <p>Use the <b>cfm md</b> command, using one of four MD naming conventions, specify an MD name to enter named MD configuration mode. See <a href="#">“MD Naming Conventions”</a> on page 27-13 for MD naming options.</p>

This example shows how to enter the system level default MD configuration command mode. Note that the command prompt changes to indicate that you have moved to Default Maintenance Domain mode for configuring default maintenance domain values.

```
7100 Chassis(rw-config)->cfm default-md default
7100 Chassis(su-config-cfm-default-md-def)->
```

This example enters default configuration mode for VLAN 20. Note that the command prompt changes to indicate that you have moved to Default Maintenance Domain mode for configuring VLAN 20.

```
7100 Chassis(rw-config)->cfm default-md vid 20
7100 Chassis(su-config-cfm-default-md.20)->
```

This example shows how to enter configuration command mode for the **myMD1** maintenance domain. This maintenance domain instance is assigned index 1.

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(rw-config-cfm-md.1)->
```

## MD Naming Conventions

When accessing the named MD configuration mode, use one of four naming conventions to identify a CFM MD:

- **String-name** – A string of up to 43 printable characters. This format provides descriptive freedom in naming the association.
- **Dns-like-name** – A string of up to 43 printable characters that represents a standard domain name server convention. This option is for management purposes. A check is done to assure that you have entered a legally formatted DNS name.
- **Mac-int-name** – A MAC address format plus an integer index ID value. You must follow a supported MAC address format as described in the parameter table.
- **No-name** – A no name option that sets a NULL string as the MD name. It can only be used for a single MD on each device.

A non-configurable index value is associated with every MD and MA. This index value appears in the prompt to provide context to the prompt. For example, the prompt (rw-config-cfm-md.1) is for the MD index 1 configuration context. Use the **show cfm md** command to display the index value for each configured MD.

## Setting SenderID TLV Permission

You configure the ID permission setting for the content sent in the SenderID TLV by the maintenance points. Given that CFM PDUs can pass through domains owned by organizations to which you may not want topology information exposed, use the ID permission settings to configure an appropriate permission level for the monitored service. The value CFM will use depends upon the context and precedence of the values configured. ID permission can be configured in CFM command modes: Default MD, Default MD VLAN, MD, MA-Comp configuration mode. The modes here are listed in precedence from lowest to highest. Configuring ID permission in MA-Comp configuration mode takes precedence over any other CFM configuration mode for the current MA context.

Enabling ID permission includes in PDUs sent by a maintenance point (MIP or MEP) informational TLVs that identify the bridge by chassis, remote management or both.

You can specify:

- **Chassis** – Includes the chassis' MAC address in the TLV
- **Manage** – Includes the method of remote management in the TLV
- **Both** – Includes both chassis and manage information in the TLV
- **None** – No SenderID TLV is sent
- **Defer** – In a VLAN default or MA component context, the SenderID TLV behavior defers to the next highest configuration level settings.

The **defer** option is only valid in the Default Maintenance Domain VLAN service context, the (named) Maintenance Domain context, or the Maintenance Association Component context. When **defer** is used in the:

- Default MD VLAN service context, the **id-permission** value defaults to the value set for the default MD
- MA component context, the **id-permission** value defaults to the value set for the (named) Maintenance Domain within which the MA component resides

Use the **id-permission** command in the appropriate mode to specify the informational content to include in the SenderID TLV. ID permission defaults to no SenderID TLVs being sent.

This example shows how to configure the maintenance points to defer to the MD setting for ID permission for the **myMA1** maintenance association component context:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->ma-comp
7100 Chassis(su-config-cfm-macomp)->id-permission defer
7100 Chassis(su-config-cfm-macomp)->enable
7100 Chassis(su-config-cfm-macomp)->
```

## Enabling Maintenance Intermediate-Points (MIP)

If MIP creation is enabled, MIPs are auto-created on all ports assigned to the CFM service. See [“Maintenance Intermediate Point \(MIP\)”](#) on page 27-6 for a description of the role MIPs play in a CFM process.

MIP creation can be configured in CFM command modes: Default MD, Default MD VLAN, MD, MA-Comp configuration mode. The modes here are listed in precedence from lowest to highest. Configuring MIP creation in MA-Comp configuration mode takes precedence over any other CFM configuration mode for the current MA context.

You can specify:

- **Default** – MIPs are created for the MD or MA context
- **Explicit** – MIPs are created for the MD or MA context, only if a MEP exists at the next lower MD level
- **None** – MIPs are not created in the MD or MA context
- **Defer** – In a VLAN default or MA component context, the MIP creation defers to the next highest configuration level settings. The **defer** option is only valid in the Default Maintenance Domain VLAN service context, the (named) Maintenance Domain context, or the Maintenance Association Component context. When **defer** is used in the:
  - Default MD VLAN service context, MIP creation defaults to the value set for the default MD
  - MA component context, MIP creation defaults to the value set for the (named) Maintenance Domain within which the MA component resides

Use the **mhf-creation** command in the appropriate configuration mode to set whether the creation of maintenance intermediate-point half function (MHF) is allowed for the current context. MIPs are not enabled by default.

This example shows how to set the system default maintenance domain value to allow an MHF to be created only when a MEP in the next lowest MD exists:

```
7100 Chassis(rw-config)->cfm default-md default
7100 Chassis(su-config-cfm-default-md-def)->mhf-creation explicit
7100 Chassis(su-config-cfm-default-md-def)->
```

## Setting the MD Level

The Ethernet CFM service network is partitioned into maintenance levels. Each maintenance level is defined by the reach and scope of the organization which administers the network equipment. Higher maintenance levels exist at the edge of the network. Network customers typically own these higher levels. Lower maintenance levels typically reside closer to the network core, and are

usually reserved for service providers or network operators. Maintenance levels are hierarchical in nature. Higher maintenance levels encapsulate lower maintenance levels.

CFM PDUs at a given level are distributed and processed among maintenance points within that specific domain. CFM PDUs from a domain that encapsulates a lower level domain, pass transparently through the enclosed lower level domain as they transit to the remote side of the domain they belong to.



**Note:** See [Figure 27-1](#) on page 27-3 for a depiction of MD encapsulation. The customer equipment MD at level 5 encapsulates both the service provider (level 3) and network operator (level 2) MDs; the service provider MD at level 3 encapsulates the network operator MDs.

CFM PDUs from a lower level domain encapsulated by a higher level domain are restricted from exiting the domain. These CFM PDUs are filtered by maintenance points of the higher level encapsulating domain.

These rules for transport and containment ensure that domain management and administration is restricted to the responsible organization. It also ensures that information concerning either the enclosed or enclosing domain is restricted to that domain. This prevents the customer from learning details concerning the internal topology of the service provider and operators and service providers from learning details concerning the customer topology.

The maintenance level to be assigned to an MD is typically determined by the service provider in the service level agreement.

Use the **level** command in the default MD or MD configuration mode to specify the MD level.

This example shows how to configure the level to 5 for the **myMD1** maintenance domain:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->level 5
7100 Chassis(su-config-cfm-md.1)->enable
7100 Chassis(su-config-cfm-md.1)->
```

## Changing the Maintenance Domain Name

The name of an MD can be changed from within the named MD context. Use the **name** command specifying a supported MD name type and name. See [“MD Naming Conventions”](#) on page 27-13 for details on naming an MD.

This example shows how to change the name of the **myMD1** maintenance domain to **yourMD1**. Note that the index number for this maintenance domain does not change.

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->name md string-name yourMD1
7100 Chassis(su-config-cfm-md.1)->enable
```

## Configuring a Maintenance Association (MA)

Refer to [“Maintenance Association \(MA\)”](#) on page 27-4 for an MA overview discussion.

For information about...	Refer to page...
<a href="#">Accessing MA Configuration Mode</a>	<a href="#">27-16</a>
<a href="#">Enabling the Maintenance Association Configuration</a>	<a href="#">27-16</a>
<a href="#">Changing the Maintenance Association Name</a>	<a href="#">27-16</a>



For information about...	Refer to page...
<a href="#">Setting the Continuity Check Message (CCM) Interval</a>	27-17
<a href="#">Configuring the Maintenance Association MEP List</a>	27-17
<a href="#">Configuring the Maintenance Association Components</a>	27-18

## Accessing MA Configuration Mode

You access MA configuration from the MD configuration mode specifying a name for the MA. One of three naming conventions can be used to identify a CFM MA:

- **String-name** – A string of up to 45 printable characters. This format provides descriptive freedom in naming the association.
- **Vid-name** – An integer value between 0 - 4094. This format restricts the association name to the VLAN range. For management purposes, show command output will label this format as a VLAN type. Use this format when the association is directly related to fault managing a VLAN.
- **Id-name** – An integer value between 0 - 65535. This format restricts the association name to an integer range. Use this format when a sequential naming scheme is being used to manage the associations.

Use the **ma** command in MD configuration command mode to enter MA configuration mode for the named MA.

This example shows how to enter configuration command mode for the **myMD1** maintenance domain and enter MA configuration mode for the maintenance association named **myMA1**:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(rw-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(rw-config-cfm-ma.1)->
```

## Enabling the Maintenance Association Configuration

Changes made in the named MA contexts do not take affect until the the CFM MA configuration is activated, using the **enable** command. See “[Activating CFM Configuration](#)” on page 27-11 for CFM configuration activation details.

This example shows how to activate **myMA1** for the **myMD1** maintenance domain:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->enable
7100 Chassis(su-config-cfm-ma.1)->
```

## Changing the Maintenance Association Name

The name of an MA can be changed from within the named MA context. Use the **name** command specifying a supported MA naming convention and name. See the MA naming convention information in “[Accessing MA Configuration Mode](#)” on page 27-16 for details on naming an MA.

This example shows how to change the MA name for the **myMA1** maintenance association to **yourMA1**:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
```

```

7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->name string-name yourMA1
7100 Chassis(su-config-cfm-ma.1)->enable
7100 Chassis(su-config-cfm-ma.1)->

```

## Setting the Continuity Check Message (CCM) Interval

The interval between CCMs can be set to 1 second, 10 seconds, 1 minute, or 10 minutes.

All maintenance end points (MEPs) in an association must be configured for the same CCM interval. The source MEP sends a continuity check message at the interval set by this command. Should the remote end-point not be configured for the same CCM interval as the source end-point, CFM logs a configuration error and potentially triggers a defect.

If the remote end-point does not receive the continuity check message within a period of 3.5 times the configured CCM interval, an error is logged. Intermediate-points (MIPs) on the path between the sending and receiving MEPs do not actively log errors.

Use the **ccm-interval** command in MA configuration mode to change the interval between the transmission of CCMs. The CCM interval defaults to 1 second.

This example sets the interval between CCM transmissions used by all MEPs in the **myMA1** maintenance association to **10** seconds:

```

7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->ccm-interval 10sec
7100 Chassis(su-config-cfm-ma.1)->enable
7100 Chassis(su-config-cfm-ma.1)->

```

## Configuring the Maintenance Association MEP List

The MA must have knowledge of the local and remote MEP IDs for the local end-points to recognize the remote end-points. All MEPs in an association must be listed in the association MEP list and the MEP list must be enabled for MEPs to be operational. MEPs are specified as a list of MEP IDs each separated by a comma (",") or, if a range of IDs, by a dash ("-").

Use the **mep-list** command, in MA configuration mode, to specify the MEPs that are or will be present in the MA. The MEP list is disabled by default. Use the **enable** option to enable the MEP list. Use the **disable** option to take a remote MEP out of defect consideration. You might want to disable a remote MEP during pre-provisioning a network, during the period in which all end-points have not yet been set up.

This example shows how to create and enable the MEP list for MEPs **30** through **35**, **1000**, and **1005** on the **myMA1** maintenance association:

```

7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->mep-list 30-35,1000,1005 enable
7100 Chassis(su-config-cfm-ma.1)->enable
7100 Chassis(su-config-cfm-ma.1)->

```

## Configuring the Maintenance Association Components

MA component configuration makes up a subset of the MA configuration within its own configuration mode. You access MA component configuration from MA configuration mode using the **ma-comp** command.

Enter the MA component configuration mode to:

- Configure ID permission for the MA context (see [“Setting SenderID TLV Permission”](#) on page 27-13)
- Configure MHF creation for the MA context (see [“Enabling Maintenance Intermediate-Points \(MIP\)”](#) on page 27-14)
- Configure the VLAN the MA is associated with (see [Setting the Maintenance Association VLAN Service](#))

### Setting the Maintenance Association VLAN Service

The service CFM monitors gets assigned to the MA. CFM supports the monitoring of a single VLAN service. If you wish to protect multiple VLANs in your system, create an MA for each VLAN to be protected. When CFM monitors a VLAN service, the monitored VLAN must be specified in the MA component configuration mode using the **vid** command. The default VLAN value for an MA is **0** (no VLAN).

This example sets VLAN **1000** as the configured VLAN for the **myMA1** component configuration:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->ma-comp
7100 Chassis(su-config-cfm-macomp)->vid 1000
7100 Chassis(su-config-cfm-macomp)->
```

### Enabling the Maintenance Association Component Configuration

Changes made in the MA component context do not take affect until the MA component configuration is activated, using the **enable** command. See [“Activating CFM Configuration”](#) on page 27-11 for CFM configuration activation details.

This example shows how to activate **myMA1** component configuration:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->ma-comp
7100 Chassis(su-config-cfm-macomp)->enable
7100 Chassis(su-config-cfm-macomp)->
```

## Configuring a Maintenance End-Point (MEP)

Refer to [“Maintenance Point \(MP\)”](#) on page 27-5 for a Maintenance Point overview discussion that includes the MEP and MIP.

For information about...	Refer to page...
<a href="#">Accessing MEP Configuration Mode</a>	27-19
<a href="#">Configuring the MEP Bridge Port</a>	27-19

For information about...	Refer to page...
<a href="#">Configuring the MEP VLAN</a>	<a href="#">27-19</a>
<a href="#">Configuring MEP Direction</a>	<a href="#">27-20</a>
<a href="#">Setting the Lowest Priority MEP Defect Alarm</a>	<a href="#">27-20</a>
<a href="#">Enabling MEP CCMs</a>	<a href="#">27-21</a>
<a href="#">Activating the MEP State Machine and the Remote MEP</a>	<a href="#">27-22</a>
<a href="#">Modifying the MEP CCM and Linktrace 802.1p Priority</a>	<a href="#">27-22</a>
<a href="#">Enabling the Maintenance End-point Configuration</a>	<a href="#">27-23</a>

## Accessing MEP Configuration Mode

You access MEP configuration mode from the MA configuration mode using the **mep** command. When accessing the MEP configuration mode, specify a MEP ID in the range **1 - 8191**.

This example enters configuration mode for the **myMA1** MEP **1000**:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->mep 1000
7100 Chassis(su-config-cfm-mep.1000)->enable
7100 Chassis(su-config-cfm-mep.1000)->
```

## Configuring the MEP Bridge Port

The bridge port the MEP is on must be configured. Use the **port** command in MEP configuration mode to configure the MEP port.

This example shows how to set **tg.1.1** as the port the **myMA1** MEP 1000 is attached to:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->mep 1000
7100 Chassis(su-config-cfm-mep.1000)->port tg.1.1
7100 Chassis(su-config-cfm-mep.1000)->enable
```

## Configuring the MEP VLAN

The MEP VLAN is the VLAN upon which the MEP transmits its PDUs. For the current CFM service implementation a MEP can only be associated with the single service VLAN, so there is no need to specify the MEP VLAN since it defaults to the service VLAN. If you do specify it, it must agree with the service VLAN. Use the **vid** command in MEP configuration mode to specify the MEP VLAN. VLANs can have a value from **0 - 4094**. The MEP VLAN defaults to **0** (no VLAN).

This example sets VLAN **1000** as the configured VLAN for the **myMA1** MEP **1000**:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->mep 1000
7100 Chassis(su-config-cfm-mep.1000)->vid 1000
7100 Chassis(su-config-cfm-mep.1000)->enable
```

```
7100 Chassis(su-config-cfm-mep.1000)->
```

## Configuring MEP Direction

The direction a MEP faces is relative to the link and the bridge relay. A down-MEP sends CFM Ethernet frames towards and receives CFM Ethernet frames from the link. An up-MEP sends CFM Ethernet frames towards the bridge relay and receives CFM Ethernet frames from the bridge relay. The bridge relay sends and receives frames to the ports associated with the CFM service.

A down-MEP performs the following functions:

- Sends and receives CFM Ethernet frames at its MD level over the link connected to the port where the MEP is configured.
- Drops all CFM Ethernet frames at its MD level or lower that come from the bridge relay.
- Processes all CFM Ethernet frames at its MD level coming from the link.
- Drops all CFM Ethernet frames at a lower MD level coming from the link.
- Transparently forwards all CFM Ethernet frames at a higher MD level, regardless of whether they came from the link or the bridge relay.

An up-MEP performs the following functions:

- Transparently forwards all CFM Ethernet frames that have a higher MD level.
- Sends and receives CFM Ethernet frames with the same MD level through the bridge relay, but not over the link.
- Drops all CFM Ethernet frames at its MD level or lower that come from the link.
- Processes all CFM Ethernet frames at its MD level coming from the bridge relay.
- Drops all CFM Ethernet frames at a lower MD level coming from the bridge relay.

MEP direction defaults to down.

This example shows how to set the direction for the **myMA1** MEP **1000** to up:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->mep 1000
7100 Chassis(su-config-cfm-mep.1000)->direction up
7100 Chassis(su-config-cfm-mep.1000)->enable
7100 Chassis(su-config-cfm-mep.1000)->
```

## Setting the Lowest Priority MEP Defect Alarm

You can set the lowest priority defect that will generate a fault alarm syslog message using the **alarm-defect-syslog** command in the MEP configuration mode.

You can set the lowest priority defect that will generate a fault alarm trap message using the **alarm-defect-trap** command in the MEP configuration mode.

The defects reported is based upon the specified level:

- **all-def** – Specifies that the DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM defects will generate a fault alarm syslog message. See [Table 27-3](#) for defect descriptions.
- **mac-rem-err-xcom** – Specifies that the DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM defects will generate a fault alarm syslog message (default).

- **rem-err-xcon** – Specifies that the DefRemoteCCM, DefErrorCCM, and DefXconCCM defects will generate a fault alarm syslog message.
- **err-xcon** – Specifies that the DefErrorCCM, and DefXconCCM defects will generate a fault alarm syslog message.
- **xcon** – Specifies that the DefXconCCM defect will generate a fault alarm syslog message.
- **no-xcon** – Specifies that no fault alarm Syslog messages will be generated.

**Table 27-3 MEP Defect Definitions**

Defect	Description
DefRDICCM	One or more continuity check messages received by this MEP contained the RDI bit. This bit indicates that some other MEP in this MEP's MA is transmitting the RDI bit. This defect clears if continuity check messages from all remote MEPs do not have the RDI bit set. Devices set the RDI bit if they have received continuity check messages from a remote MEP that indicates the remote MEP has seen defect notifications from other MEPs in the MA.
DefMACstatus	The port status is not indicating "UP" for all remote MEPs on received continuity check messages, or the interface status for any remote MEP on received continuity check messages is not indicating "UP".
DefRemoteCCM	This MEP is not receiving continuity check messages from a MEP in its configured list.
DefErrorCCM	This MEP is receiving continuity check messages from a remote MEP that either uses an invalid MEP ID or uses a continuity check interval that does not match the receiving MEP.
DefXconCCM	This MEP is receiving continuity check messages from a remote MEP that either uses an MD level lower than the receiving MEP or uses a different MD or MA name than the receiving MEP.

You can set the minimum time a defect must be present before an alarm is generated using the **alarm-time** command. You can set the minimum amount of time a defect must be absent before and alarm is reset using the **reset-time** command. Alarm time and reset time values are set in centiseconds between **250** (2.5 seconds) and **1000** (10 seconds). The default value is **2.5** seconds.

This example sets the Syslog defect alarm setting for the **myMA1** MEP **1000** to generate a Syslog message if any supported defect is present for at least 3 seconds:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->mep 1000
7100 Chassis(su-config-cfm-mep.1000)->alarm-defect-syslog all-def
7100 Chassis(su-config-cfm-mep.1000)->alarm-time 300
7100 Chassis(su-config-cfm-mep.1000)->reset-time 300
7100 Chassis(su-config-cfm-mep.1000)->enable
```

## Enabling MEP CCMs

By default a MEP does not send CCMs. Use the **cci-enabled** command in MEP configuration mode to enable the sending of CCMs for this MEP.

This example shows how to enable the sending of continuity check messages for MEP **1000** on **myMA1**:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
```

```
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->mep 1000
7100 Chassis(su-config-cfm-mep.1000)->cci-enabled
7100 Chassis(su-config-cfm-mep.1000)->enable
7100 Chassis(su-config-cfm-mep.1000)->
```

## Activating the MEP State Machine and the Remote MEP

The MEP state machine must be set to active for the MEP to be operational using the **active** command. If the MEP state machines are not active, the MEP will not operate.

If the state machine status is active and the MEP configuration is not enabled using the **enable** command, the MEP will have a Row Status of "Active" in the **show cfm md ma mep remote-mep** command, but will not perform any functions such as send and receive PDUs or generate defects.

The remote MEP(s) for this service must be enabled in the database for the CFM service to be operational between the two MEPs. Use the **remote-mep active** command to enable any remote MEPs associated with the local MEP. If you want to remove a remote MEP from defect consideration, for example, while provisioning the network, you can disable the remote MEP using the **no remote-mep active** command.

This example sets the administrative state for the **myMA1 MEP 1000** to active and enables remote MEP **2000** in the MEP database:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->mep 1000
7100 Chassis(su-config-cfm-mep.1000)->active
7100 Chassis(su-config-cfm-mep.1000)->remote-mep active 2000
7100 Chassis(su-config-cfm-mep.1000)->enable
7100 Chassis(su-config-cfm-mep.1000)->
```

## Modifying the MEP CCM and Linktrace 802.1p Priority

You can modify the 802.1p priority used by connectivity check messages and messages sent by the linktrace protocol. 802.1p CCM and linktrace message priority is on a per MEP basis. Valid values are 0 - 7. The default value is 1. Use the **priority** command in MEP configuration mode to modify the 802.1p priority used by CCMs and linktrace messages.

The **priority** command does not affect loopback protocol message 802.1p priority. Loopback protocol message priority can be modified when entering the **loopback** command.

This example shows how to set set the continuity check message and linkstate message priority for MEP **1000** to **3**:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->mep 1000
7100 Chassis(su-config-cfm-mep.1000)->priority 3
7100 Chassis(su-config-cfm-mep.1000)->
```

## Enabling the Maintenance End-point Configuration

Changes made in the MEP context do not take effect until the MEP configuration is activated, using the **enable** command. See [“Activating CFM Configuration”](#) on page 27-11 for CFM configuration activation details.

This example shows how to enable MEP 1000 on **myMA1**:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->mep 1000
7100 Chassis(su-config-cfm-mep.1000)->enable
7100 Chassis(su-config-cfm-mep.1000)->
```

## CFM Loopback and Linktrace Protocols

There are two CFM diagnostic protocols available in MEP configuration mode: loopback and linktrace.

### The CFM Loopback Protocol

The loopback protocol sends loopback messages (LBM) to either a specified maintenance end-point (MEP) or to an MP MAC address. The CFM loopback protocol is similar to IP ICMP ping. The CFM loopback protocol displays whether there is connectivity between the initiating device and the target device. If an operational path to the MEP or MIP exists, the remote MEP or MIP will respond. If no response is received by the source MEP, no operational path exists. If no response is received use the linktrace protocol ([“The CFM Linktrace Protocol”](#) on page 27-23) to help verify where in the path the problem occurred.

If you specify a MEP, the specified remote MEP must have already communicated with this MEP and there must be an entry in the MEP database for the remote MEP. The loopback will fail if the remote MEP is specified and the specified MEP has not yet communicated with this MEP.

The LBM 802.1p priority can be configured when initiating a loopback.

This example shows how to send 5 loopback messages to MAC address **01:3a:b2:af:65:de** from the **myMA1 MEP 1000**:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->mep 1000
7100 Chassis(su-config-cfm-mep.1000)->loopback mac 01:3a:b2:af:65:de messages 5
Sending 5 Ethernet CFM loopback messages to 01-3a-b2-af-65-de
...
Success rate is 100 percent (5/5)
7100 Chassis(su-config-cfm-mep.1000)->
```

### The CFM Linktrace Protocol

The CFM linktrace protocol is used to help verify a path and identify where in a path a connectivity problem exists by indicating that an incomplete path between the initiating device and the target device exists. The CFM linktrace protocol is similar to IP traceroute. Linktrace messages (LTM) are sent to either a specified MEP or to an MP MAC address. If you specify a MEP, the specified remote MEP must have already communicated with this MEP and there must



be an entry in the MEP database for the remote MEP. The linktrace will fail if the remote MEP is specified and the specified MEP has not yet communicated with this MEP.

You can specify a TTL which sets the maximum number of network hops the LTM will traverse before expiring.

You can specify whether each hop along the linktrace path uses the local filter database exclusively or is allowed to also use the local MIP CCM database to determine reachability to the target.

The LTM shares the same priority as the continuity check message and can not be separately configured. The priority for continuity check messages and LTMs is set using the **priority** command in MEP configuration mode.

This example shows how to send linktrace messages to MEP 2000 from the **myMA1** MEP 1000:

```
7100 Chassis(rw-config)->cfm md string-name myMD1
7100 Chassis(su-config-cfm-md.1)->ma string-name myMA1
7100 Chassis(su-config-cfm-ma.1)->mep 1000
7100 Chassis(su-config-cfm-mep.1000)->linktrace mep 2000
```

Linktrace to 00-00-00-10-00-03, Transaction ID 29481

MD Name: abc

MA Name: abc

MEP ID : 1, Interface ge.4.18

```
=====
Hop  TTL      Source MAC      Next hop MAC      Relay
-----
  1   63  00-1f-45-9e-3e-d1  00-00-00-10-00-00  MIP-DB
  2   62  00-00-00-10-00-00  00-00-00-00-00-00  Hit
```

```
7100 Chassis(su-config-cfm-mep.1000)->
```

## Configuring Connectivity Fault Management

[Table 27-4](#) lists the 7100-Series device default CFM configuration settings.

**Table 27-4 Default Connectivity Fault Management Configuration Settings**

Parameter	Description	Default Value
CFM global state	The global state of CFM as configured in global configuration mode.	Disabled
CFM logging filter	The ability to limit the sending of CFM Syslog messages by MD, MA, and MEP.	All Syslog messages are sent
ID permission	The content sent in the SenderID TLV by maintenance points.	None
MD level	The maintenance domain level.	0
MHF creation	Sets whether the creation of MIP half functions is allowed.	None

**Table 27-4 Default Connectivity Fault Management Configuration Settings (continued)**

Parameter	Description	Default Value
CFM activation	The activation of CFM configuratin changes in an MD, MA, MA component, or MEP configuration context.	Not activated
CCM interval	The interval between the transmission of continuity check messages (CCM).	1 second
VLAN association	The VLAN the MA and MEP are associated with.	0
MEP state machine admin state	The administrative state of the MEP state machine.	Inactive
Minimum alarm defect reported	The minimum defect that will cause a Syslog or trap to be sent.	mac-rem-err-xcom
Alarm time	The minimum time a defect must be present before an alarm is generated.	2.5 seconds
CCI enabled	Determines whether continuity check messages (CCM) are enabled for generation and reception by a MEP.	Disabled
MEP direction	The direction a MEP faces which determines whether CFM PDUs are sent towards the link or bridge relay.	Down
CCM 802.1p priority	The 802.1p priority setting for CCMs and linktrace messages.	1
Remote MEP active	Determines the remote MEP state for the MEP being configured.	Disabled
Defect reset time	The time a MEP defect must be absent before an alarm is reset.	2.5 seconds

[Table 27-5](#) lists CFM global configuration commands. Global configuration commands provide access to the CFM global default MD and monitored VLAN service configuration modes and the MD configuration mode (for a specific MD), the global enabling of CFM, and CFM logging filter configuration. These commands are accessed in global configuration mode.

**Table 27-5 CFM Global Configuration**

Task	Command
Access the CFM system default configuration mode.	<b>cfm default-md default</b>
Access the default configuration mode for a monitored VLAN service.	<b>cfm default-md vid <i>vlan-id</i></b>
Globally enable CFM on the device. CFM is disabled by default	<b>cfm enable</b>
Filter the sending of CFM Syslog messages by maintenance domains (MDs), maintenance associations (MAs), and maintenance end points (MEPs). All CFM Syslog messages are sent by default.	<b>cfm logging filter md {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i>} [ma string-name <i>name</i>   vid-name <i>vid-name</i>   id-name <i>id-name</i>   index <i>index</i>] [mep <i>mep-id</i>]</b>
Enter Maintenance Domain (MD) Configuration mode for a specific named MD. If the maintenance domain does not exist, this command will create it.	<b>cfm md {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name}</b>

**Table 27-5 CFM Global Configuration (continued)**

Task	Command
Optionally, configure a VLAN table to associate one or more CFM services with a primary CFM service	<b>cfm vlan-table primary-selector</b> <i>primary-selector selector-list selector-list</i> [ <b>enable</b>   <b>disable</b> ]

[Table 27-6](#) on page 27-26 lists CFM system MD and monitored VLAN service default configuration commands. Enter the appropriate default MD configuration mode using commands listed in [Table 27-5](#). In CFM system default configuration mode, entered default configuration is applied globally. In monitored VLAN service default configuration mode, entered default configuration is applied to the specified VLAN.

**Table 27-6 CFM Global and Monitored VLAN Service Default Configuration**

Task	Command
Optionally, configure the ID permission setting, within the appropriate default configuration mode, for the content sent in the SenderID TLV by the maintenance points. Defaults to <b>none</b> .  The <b>defer</b> option is not supported in the global MD default command mode.	<b>id-permission</b> { <b>chassis</b>   <b>manage</b>   <b>chassis-manage</b>   <b>none</b>   <b>defer</b> }
Optionally, set the default maintenance domain level for the current default context. The MD level defaults to <b>0</b> .	<b>level</b> <i>level</i>
Optionally, set whether the creation of maintenance intermediate-point half function (MHF) is allowed for the current default context. Defaults to <b>none</b> .	<b>mhf-creation</b> { <b>default</b>   <b>explicit</b>   <b>none</b> }

[Procedure 27-1](#) describes how to configure a MD. All commands listed in this procedure are entered in the MD configuration mode. See [Table 27-5](#) on page 27-25 for the command to access the MD configuration mode.

**Procedure 27-1 CFM Maintenance Domain (MD) Configuration**

Step	Task	Command(s)
1.	Optionally, change the MD name for the current MD context.	<b>name</b> { <b>string-name</b> <i>name</i>   <b>dns-like-name</b> <i>dns-name</i>   <b>mac-int-name</b> <i>mac-name</i>   <b>no-name</b> }
2.	Optionally, configure the ID permission setting for the content sent in the SenderID TLV by the maintenance points. Defaults to <b>none</b> .	<b>id-permission</b> { <b>chassis</b>   <b>manage</b>   <b>chassis-manage</b>   <b>none</b>   <b>defer</b> }
3.	Optionally, set the maintenance domain level. The MD level defaults to <b>0</b> .	<b>level</b> <i>level</i>
4.	Optionally, set whether the creation of maintenance intermediate-point half function (MHF) is allowed. Defaults to <b>none</b> .	<b>mhf-creation</b> { <b>default</b>   <b>explicit</b>   <b>none</b>   <b>defer</b> }
5.	Activate the CFM configuration for the named MD context.	<b>enable</b>
6.	Optionally, enter Maintenance Association (MA) Configuration mode for the specified MA.	<b>ma</b> { <b>string-name</b> <i>name</i>   <b>vid-name</b> <i>vlan</i>   <b>id-name</b> <i>id</i> }

[Procedure 27-2](#) describes how to configure an MA. All commands listed in this procedure are entered in the MA configuration mode. See [Procedure 27-1](#) for the command to access the MA configuration mode.

### Procedure 27-2 CFM Maintenance Association (MA) Configuration

Step	Task	Command(s)
1.	Specify the MEPs that are or will be present in the MA.	<b>mep-list</b> <i>mep-list</i> [ <b>enable</b>   <b>disable</b> ]
2.	Optionally, set the interval between continuity check messages (CCM)s. Defaults to 1 seconds.	<b>ccm-interval</b> { <b>1sec</b>   <b>10sec</b>   <b>1min</b>   <b>10min</b> }
3.	Optionally, change the Maintenance Association name for the current MA context.	<b>name</b> { <b>string-name</b> <i>name</i>   <b>vid-name</b> <i>vid-name</i>   <b>id-name</b> <i>id-name</i> }
4.	Optionally, enter MA component configuration mode for the MA.	<b>ma-comp</b>
5.	Activate the CFM configuration for the current MA context.	<b>enable</b>
6.	Optionally, enter Maintenance End-Point (MEP) Configuration mode for the specified end-point.	<b>mep</b> <i>mep-id</i>

[Procedure 27-3](#) describes how to configure the MA components for the current MA context. All commands listed in this procedure are entered in the MA-Comp configuration mode. See [Procedure 27-2](#) for the command to access the MA-Comp configuration mode.

### Procedure 27-3 CFM Maintenance Association Component (MA-Comp) Configuration

Step	Task	Command(s)
1.	Specify the VLAN the maintenance association is associated with	<b>vid</b> <i>vlan-id</i>
2.	Optionally, configure the ID permission setting for the content sent in the SenderID TLV by the maintenance points for this MA context. Defaults to <b>none</b> .	<b>id-permission</b> { <b>chassis</b>   <b>manage</b>   <b>chassis-manage</b>   <b>none</b>   <b>defer</b> }
3.	Optionally, set whether the creation of maintenance intermediate-point half function (MHF) is allowed. Defaults to <b>none</b> .	<b>mhf-creation</b> { <b>default</b>   <b>explicit</b>   <b>none</b> }
4.	Activate the CFM configuration for the named MD context.	<b>enable</b>

[Procedure 27-4](#) describes how to configure a MEP for the current MA context. All commands listed in this procedure are entered in the MEP configuration mode. See [Procedure 27-2](#) on page 27-27 for the command to access the MEP configuration mode.

**Procedure 27-4 CFM Maintenance Association End-Point (MEP) Configuration**

Step	Task	Command(s)
1.	Optionally, set the lowest priority defect that will generate a fault alarm syslog message. Default is <b>mac-rem-err-xcon</b> .	<b>alarm-defect-syslog</b> { <b>all-def</b>   <b>mac-rem-err-xcon</b>   <b>rem-err-xcon</b>   <b>err-xcon</b>   <b>xcon</b>   <b>no-xcon</b> }
2.	Optionally, set the lowest priority defect that will generate a fault alarm trap message. Default is <b>mac-rem-err-xcon</b> .	<b>alarm-defect-trap</b> { <b>all-def</b>   <b>mac-rem-err-xcon</b>   <b>rem-err-xcon</b>   <b>err-xcon</b>   <b>xcon</b>   <b>no-xcon</b> }
3.	Optionally, set the minimum time a defect must be present before an alarm is generated. Default is <b>2.5</b> seconds.	<b>alarm-time</b> <i>time</i>
4.	Optionally, enable generation of continuity check messages. Default is disabled.	<b>cci-enabled</b>
5.	Optionally, configure whether the MEP faces the bridge relay (up) or the bridge port (down). Default is <b>down</b>	<b>direction</b> { <b>down</b>   <b>up</b> }
6.	Optionally, transmit CFM linktrace messages to the specified MEP or MAC address to help verify a path and identify where in a path a connectivity problem exists.	<b>tracelink</b> { <b>mep</b> <i>mep-id</i>   <b>mac</b> <i>mac-addr</i> } [ <b>tll</b> <i>time-to-live</i> ] [ <b>fdb-only</b> ]
7.	Optionally, transmit CFM loopback messages to the specified MEP or MAC address to verify an operational path exists to the MEP.	<b>loopback</b> { <b>mep</b> <i>mep-id</i>   <b>mac</b> <i>mac-addr</i> } [ <b>messages</b> <i>num-messages</i> ] [ <b>priority</b> <i>priority</i> ] [ <b>data</b> <i>data</i> ]
8.	Configure the bridge port the MEP is attached to.	<b>port</b> <i>port</i>
9.	Optionally, configure the 802.1 priority for continuity check messages and linktrace message sent by this MEP. Default is <b>1</b> .	<b>priority</b> <i>priority</i>
10.	Enable a remote MEP in the database for the current MEP.	<b>remote-mep</b> <i>remote-mep-id</i> <b>active</b>
11.	Optionally, configure the time a MEP defect must be absent before an alarm is reset. Default is <b>2.5</b> seconds.	<b>reset-time</b> <i>time</i>
12.	Activate the administrative state of the MEP (maintenance end point) state machine. Defaults to inactive.	<b>active</b>
13.	Enable the CFM configuration for current MEP context. Defaults to disabled	<b>enable</b>

[Table 27-7](#) describes how to manage CFM. CFM clear commands can be entered in any command mode.

**Table 27-7 CFM Management Commands**

Task	Command(s)
To clear the CFM bridge MIP CCM database.	<b>clear cfm bridge mip-ccm</b>

**Table 27-7 CFM Management Commands (continued)**

Task	Command(s)
To clear the check continuity message database for a specified MEP or all end-points for the specified context.	<p>For a specified MD:</p> <pre><b>clear cfm ccm-database md</b> {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i>} [mep <i>mep-id</i>]</pre> <p>For the specified or all endpoints in the specified MA:</p> <pre><b>clear cfm ccm-database md</b> {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i>} <b>ma</b> {string-name <i>name</i>   vid-name <i>vlan</i>   id-name <i>id</i>   index <i>index</i>} [mep <i>mep-id</i>]</pre>
To clear the MEP counters for a specified MEP or all end-points for the specified context.	<p>For the specified or all end-points in the specified MD:</p> <pre><b>clear cfm counters md</b> {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i>} [mep <i>mep-id</i>]</pre> <p>For the specified or all endpoints in the specified MA:</p> <pre><b>clear cfm counters md</b> {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i>} <b>ma</b> {string-name <i>name</i>   vid-name <i>vlan</i>   id-name <i>id</i>   index <i>index</i>} [mep <i>mep-id</i>]</pre> <p>For the specified end-point:</p> <pre><b>clear cfm counters md</b> {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i>} <b>ma</b> {string-name <i>name</i>   vid-name <i>vlan</i>   id-name <i>id</i>   index <i>index</i>} <b>mep</b> <i>mep-id</i></pre>

[Table 27-8](#) describes how to display CFM configuration and statistics. CFM show commands can be entered in any command mode.

**Table 27-8 CFM Show Commands**

Task	Command(s)
To display CFM status and configuration for all CFM MDs, associations, and end-points.	<b>show cfm all</b>
To display maintenance intermediate point (MIP) continuity check message database entries for MHFs that do not belong to a specific MD and MA.	<b>show cfm bridge mip-ccm</b> [vid <i>vlan-id</i> ]
To display the system level default MD or CFM service default values.	<b>show cfm default-md</b> [default   vid <i>vlan-id</i> ]
To display MD information for all or the specified MD.	<b>show cfm md</b> [string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i> ]

**Table 27-8 CFM Show Commands (continued)**

Task	Command(s)
To display configuration and status for a specified or all MAs for the specified MD.	<b>show cfm md</b> {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i> } <b>ma</b> [string-name <i>name</i>   vid-name <i>vlan</i>   id-name <i>id</i>   index <i>index</i> ]
To display the MA component configuration information.	<b>show cfm md</b> {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i> } <b>ma</b> {string-name <i>name</i>   vid-name <i>vlan</i>   id-name <i>id</i>   index <i>index</i> } <b>ma-comp</b>
To display the MEP configuration information.	<b>show cfm md</b> {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i> } <b>ma</b> {string-name <i>name</i>   vid-name <i>vlan</i>   id-name <i>id</i>   index <i>index</i> } <b>mep</b> [mep-id <i>mep-id</i> ] [-verbose]
To display the error conditions in the MEP continuity check message database.	<b>show cfm md</b> {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i> } <b>ma</b> {string-name <i>name</i>   vid-name <i>vlan</i>   id-name <i>id</i>   index <i>index</i> } <b>mep</b> [mep-id <i>mep-id</i> ] <b>ccm-errors</b> [-verbose]
To display linktrace database information for all or the specified end-point.	<b>show cfm md</b> {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i> } <b>ma</b> {string-name <i>name</i>   vid-name <i>vlan</i>   id-name <i>id</i>   index <i>index</i> } <b>mep</b> [mep-id <i>mep-id</i> ] <b>linktrace</b> [-verbose]
to display the MA MEP list.	<b>show cfm md</b> {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i> } <b>ma</b> {string-name <i>name</i>   vid-name <i>vlan</i>   id-name <i>id</i>   index <i>index</i> } <b>mep-list</b>
to display the MEP's remote MEP configuration information.	<b>show cfm md</b> {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i> } <b>ma</b> {string-name <i>name</i>   vid-name <i>vlan</i>   id-name <i>id</i>   index <i>index</i> } <b>mep</b> [mep-id <i>mep-id</i> ] <b>remote-mep</b> [mep-id <i>mep-id</i> ] [-verbose]

**Table 27-8 CFM Show Commands (continued)**

Task	Command(s)
to display maintenance intermediate point (MIP) continuity check message information by MD, MA, or MEP.	Display by MD: <b>show cfm md</b> {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i> } mip-ccm [vid <i>vlan-id</i> ] Display by MA: <b>show cfm md</b> {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i> } ma {string-name <i>name</i>   vid-name <i>vlan</i>   id-name <i>id</i>   index <i>index</i> } mip-ccm [vid <i>vlan-id</i> ] Display by MEP: <b>show cfm md</b> {string-name <i>name</i>   dns-like-name <i>dns-name</i>   mac-int-name <i>mac-name</i>   no-name   index <i>index</i> } ma {string-name <i>name</i>   vid-name <i>vlan</i>   id-name <i>id</i>   index <i>index</i> } mep [mep-id <i>mep-id</i> ] mip-ccm [vid <i>vlan-id</i> ]
To display the stack table which contains maintenance point information for each maintenance point for the device (MEP and MIP).	<b>show cfm stack-table</b>
To display the global CFM status for this device.	<b>show cfm status</b>


Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

## Single MD Configuration Example

A CFM monitored service can be configured within a single MD if you have administrative control of all the devices the service passes through. For our single MD CFM example we will monitor two VLANs. See [Figure 27-6](#) for an overview of the CFM single MD configuration example.



Figure 27-6 Single MD Example Configuration Overview



The MD is named `mdCE1` and uses the string-name naming convention. An MA is required for each monitored service. The MA monitoring VLAN 100 is named `maCE1`. The MA monitoring VLAN 200 is named `maCE2`. Each MA name uses the string-name naming convention. The naming format for each MEP enabled device is the name of the MA it belongs to with a `:x` where `x` is a numeric value. Each MA has a switch internal to the MD through which the monitored service passes that will be MIP enabled. Since there is only one MD to protect, and you are not concerned that CCMs will expose the network topology to any other operators, all MEPs in the example are down MEPs. Each MA has three MEPs configured.

## Configuring Device `maCE1:1`

To configure `maCE1:1`:

- Enable CFM in global configuration mode
- Limit logging to MA `maCE1`
- Access MD mode for MD `mdCE1` using the string-name naming convention
- Set the MD level to 5
- Enable the MD configuration
- Access MA mode for MA `maCE1` using the string-name naming convention
- Use the default CCM message interval of 1 second (not configured)
- Configure the MEP list for the two down MEPs: 101, 102, and 103

- Enable the MA configuration
- Access the MA-Comp mode for **maCE1**
- Set VLAN **100** as the monitored service
- Enable MA-Comp configuration for **maCE1**
- Access MEP configuration for MEP **101**
- Set the MEP port to **tg.1.1**
- Set the MEP VLAN to **100**
- Set the MEP direction to **down** (optional because down is the direction default)
- Enable the sending of CCM messages for MEP **101**
- Set the remote MEP that CCM messages will be sent to **102** and **103**
- Activate the MEP state machines
- Enable the MEP configuration

### Device maCE1:1 CLI Input

```

maCE1:1(rw-config)->cfm enable
maCE1:1(rw-config)->cfm logging filter md string-name mdCE1 ma string-name maCE1
maCE1:1(rw-config)->cfm md string-name mdCE1
maCE1:1(rw-config-cfm-md.1)->level 5
maCE1:1(rw-config-cfm-md.1)->enable
maCE1:1(rw-config-cfm-md.1)->ma string-name maCE1
maCE1:1(rw-config-cfm-ma.1)->mep-list 101-103 enable
maCE1:1(rw-config-cfm-ma.1)->enable
maCE1:1(rw-config-cfm-ma.1)->ma-comp
maCE1:1(rw-config-cfm-macomp)->vid 100
maCE1:1(rw-config-cfm-macomp)->enable
maCE1:1(rw-config-cfm-macomp)->exit
maCE1:1(rw-config-cfm-ma.1)->mep 101
maCE1:1(rw-config-cfm-mep.101)->port tg.1.1
maCE1:1(rw-config-cfm-mep.101)->vid 100
maCE1:1(rw-config-cfm-mep.101)->direction down
maCE1:1(rw-config-cfm-mep.101)->cci-enabled
maCE1:1(rw-config-cfm-mep.101)->remote-mep active 102,103
maCE1:1(rw-config-cfm-mep.101)->active
maCE1:1(rw-config-cfm-mep.101)->enable
maCE1:1(rw-config-cfm-mep.101)->exit
maCE1:1(rw-config-cfm-ma.1)->exit
maCE1:1(rw-config-cfm-ma.1)->exit
maCE1:1(rw-config)->

```

### Configuring Device maCE1:2

To configure maCE1:2 use the information listed in [“Configuring Device maCE1:1”](#) on page 27-32 until you get to MEP configuration. At MEP configuration proceed here:

- Access MEP configuration for MEP **102**
- Set the MEP port to **tg.1.2**
- Set the MEP VLAN to **100**
- Set the MEP direction to **down** (optional because down is the direction default)
- Enable the sending of CCM messages for MEP **102**
- Set the remote MEP that CCM messages will be sent to **101** and **103**
- Activate the MEP state machines
- Enable the MEP configuration

### Device maCE1:2 CLI Input

```
maCE1:2 (rw-config) -> cfm enable
maCE1:2 (rw-config) -> cfm logging filter md string-name mdCE1 ma string-name maCE1
maCE1:2 (rw-config) -> cfm md string-name mdCE1
maCE1:2 (rw-config-cfm-md.1) -> level 5
maCE1:2 (rw-config-cfm-md.1) -> enable
maCE1:2 (rw-config-cfm-md.1) -> ma string-name maCE1
maCE1:2 (rw-config-cfm-ma.1) -> mep-list 101-103 enable
maCE1:2 (rw-config-cfm-ma.1) -> enable
maCE1:2 (rw-config-cfm-ma.1) -> ma-comp
maCE1:2 (rw-config-cfm-macomp) -> vid 100
maCE1:2 (rw-config-cfm-macomp) -> enable
maCE1:2 (rw-config-cfm-macomp) -> exit
maCE1:2 (rw-config-cfm-ma.1) -> mep 101
maCE1:2 (rw-config-cfm-mep.102) -> port tg.1.2
maCE1:2 (rw-config-cfm-mep.102) -> vid 100
maCE1:2 (rw-config-cfm-mep.102) -> direction down
maCE1:2 (rw-config-cfm-mep.102) -> cci-enabled
maCE1:2 (rw-config-cfm-mep.102) -> remote-mep active 101,103
maCE1:2 (rw-config-cfm-mep.102) -> active
maCE1:2 (rw-config-cfm-mep.102) -> enable
maCE1:2 (rw-config-cfm-mep.102) -> exit
maCE1:2 (rw-config-cfm-ma.1) -> exit
maCE1:2 (rw-config-cfm-ma.1) -> exit
maCE1:2 (rw-config) ->
```

### Configuring Device maCE1:3

To configure device maCE1:3 use the information listed in [“Configuring Device maCE1:1”](#) on page 27-32 until you get to MEP configuration. At MEP configuration proceed here:

- Access MEP configuration for MEP **103**
- Set the MEP port to **tg.1.3**
- Set the MEP VLAN to **100**
- Set the MEP direction to **down** (optional because down is the direction default)

- Enable the sending of CCM messages for MEP **103**
- Set the remote MEP that CCM messages will be sent to **101** and **102**
- Activate the MEP state machines
- Enable the MEP configuration

### Device maCE1:3 CLI Input

```

maCE1:3(rw-config)->cfm enable
maCE1:3(rw-config)->cfm logging filter md string-name mdCE1 ma string-name maCE1
maCE1:3(rw-config)->cfm md string-name mdCE1
maCE1:3(rw-config-cfm-md.1)->level 5
maCE1:3(rw-config-cfm-md.1)->enable
maCE1:3(rw-config-cfm-md.1)->ma string-name maCE1
maCE1:3(rw-config-cfm-ma.1)->mep-list 101-103 enable
maCE1:3(rw-config-cfm-ma.1)->enable
maCE1:3(rw-config-cfm-ma.1)->ma-comp
maCE1:3(rw-config-cfm-macomp)->vid 100
maCE1:3(rw-config-cfm-macomp)->enable
maCE1:3(rw-config-cfm-macomp)->exit
maCE1:3(rw-config-cfm-ma.1)->mep 101
maCE1:3(rw-config-cfm-mep.103)->port tg.1.2
maCE1:3(rw-config-cfm-mep.103)->vid 100
maCE1:3(rw-config-cfm-mep.103)->direction down
maCE1:3(rw-config-cfm-mep.103)->cci-enabled
maCE1:3(rw-config-cfm-mep.103)->remote-mep active 101,103
maCE1:3(rw-config-cfm-mep.103)->active
maCE1:3(rw-config-cfm-mep.103)->enable
maCE1:3(rw-config-cfm-mep.103)->exit
maCE1:3(rw-config-cfm-ma.1)->exit
maCE1:3(rw-config-cfm-ma.1)->exit
maCE1:3(rw-config)->

```

## Configuring Switch 1

To configure CE Device 1:

- Enable CFM in global configuration mode
- Limit logging to MA **maCE1**
- Configure the default MD VLAN to **100**
- Configure the default MD VLAN 100 MD level to **5**
- Set the default MD VLAN 100 MHF create setting to **default**

### Switch 1 CLI Input

```

Switch 1(rw-config)->cfm enable
Switch 1(rw-config)->cfm logging filter md string-name mdCE1 ma string-name maCE1
Switch 1(rw-config)->cfm default-md vid 100

```

```
Switch 1 (rw-config-def-md.100)->level 5
Switch 1 (rw-config-def-md.100)->mhf-creation default
Switch 1 (rw-config-def-md.100)->exit
Switch 1 (rw-config)->
```

## Configuring Device maCE2:1

To configure maCE2:1:

- Enable CFM in global configuration mode
- Limit logging to MA **maCE2**
- Access MD mode for MD **mdCE1** using the string-name naming convention
- Set the MD level to 5
- Enable the MD configuration
- Access MA mode for MA **maCE2** using the string-name naming convention
- Use the default CCM message interval of 1 second (not configured)
- Configure the MEP list for the two down MEPs: **201, 202, and 203**
- Enable the MA configuration
- Access the MA-Comp mode for **maCE2**
- Set VLAN **200** as the monitored service
- Enable MA-Comp configuration for **maCE2**
- Access MEP configuration for MEP **201**
- Set the MEP port to **tg.1.4**
- Set the MEP VLAN to **200**
- Set the MEP direction to **down** (optional because down is the direction default)
- Enable the sending of CCM messages for MEP **201**
- Set the remote MEP that CCM messages will be sent to **202 and 203**
- Activate the MEP state machines
- Enable the MEP configuration

### Device maCE2:1 CLI Input

```
maCE2:1 (rw-config)->cfm enable
maCE2:1 (rw-config)->cfm logging filter md string-name mdCE1 ma string-name maCE2
maCE2:1 (rw-config)->cfm md string-name mdCE1
maCE2:1 (rw-config-cfm-md.1)->level 5
maCE2:1 (rw-config-cfm-md.1)->enable
maCE2:1 (rw-config-cfm-md.1)->ma string-name maCE2
maCE2:1 (rw-config-cfm-ma.2)->mep-list 201-203 enable
maCE2:1 (rw-config-cfm-ma.2)->enable
maCE2:1 (rw-config-cfm-ma.2)->ma-comp
maCE2:1 (rw-config-cfm-macomp)->vid 200
maCE2:1 (rw-config-cfm-macomp)->enable
```

```

maCE2:1 (rw-config-cfm-macomp) ->exit
maCE2:1 (rw-config-cfm-ma.2) ->mep 201
maCE2:1 (rw-config-cfm-mep.201) ->port tg.1.4
maCE2:1 (rw-config-cfm-mep.201) ->vid 200
maCE2:1 (rw-config-cfm-mep.201) ->direction down
maCE2:1 (rw-config-cfm-mep.201) ->cci-enabled
maCE2:1 (rw-config-cfm-mep.201) ->remote-mep active 202,203
maCE2:1 (rw-config-cfm-mep.201) ->active
maCE2:1 (rw-config-cfm-mep.201) ->enable
maCE2:1 (rw-config-cfm-mep.201) ->exit
maCE2:1 (rw-config-cfm-ma.2) ->exit
maCE2:1 (rw-config-cfm-ma.2) ->exit
maCE2:1 (rw-config) ->

```

## Configuring Device maCE2:2

To configure maCE2:2 use the information listed in “[Configuring Device maCE1:1](#)” on page 27-32 until you get to MEP configuration. At MEP configuration proceed here:

- Access MEP configuration for MEP **202**
- Set the MEP port to **tg.1.5**
- Set the MEP VLAN to **200**
- Set the MEP direction to **down** (optional because down is the direction default)
- Enable the sending of CCM messages for MEP **202**
- Set the remote MEP that CCM messages will be sent to **201** and **203**
- Activate the MEP state machines
- Enable the MEP configuration

### Device maCE2:2 CLI Input

```

maCE2:2 (rw-config) ->cfm enable
maCE2:2 (rw-config) ->cfm logging filter md string-name mdCE1 ma string-name maCE1
maCE2:2 (rw-config) ->cfm md string-name mdCE1
maCE2:2 (rw-config-cfm-md.1) ->level 5
maCE2:2 (rw-config-cfm-md.1) ->enable
maCE2:2 (rw-config-cfm-md.1) ->ma string-name maCE1
maCE2:2 (rw-config-cfm-ma.2) ->mep-list 101-103 enable
maCE2:2 (rw-config-cfm-ma.2) ->enable
maCE2:2 (rw-config-cfm-ma.2) ->ma-comp
maCE2:2 (rw-config-cfm-macomp) ->vid 100
maCE2:2 (rw-config-cfm-macomp) ->enable
maCE2:2 (rw-config-cfm-macomp) ->exit
maCE2:2 (rw-config-cfm-ma.2) ->mep 101
maCE2:2 (rw-config-cfm-mep.202) ->port tg.1.2
maCE2:2 (rw-config-cfm-mep.202) ->vid 100
maCE2:2 (rw-config-cfm-mep.202) ->direction down

```

```

maCE2:2 (rw-config-cfm-mep.202) ->cci-enabled
maCE2:2 (rw-config-cfm-mep.202) ->remote-mep active 101,103
maCE2:2 (rw-config-cfm-mep.202) ->active
maCE2:2 (rw-config-cfm-mep.202) ->enable
maCE2:2 (rw-config-cfm-mep.202) ->exit
maCE2:2 (rw-config-cfm-ma.2) ->exit
maCE2:2 (rw-config-cfm-ma.2) ->exit
maCE2:2 (rw-config) ->

```

## Configuring Device maCE2:3

To configure device maCE1:3 use the information listed in [“Configuring Device maCE1:1”](#) on page 27-32 until you get to MEP configuration. At MEP configuration proceed here:

- Access MEP configuration for MEP **203**
- Set the MEP port to **tg.1.6**
- Set the MEP VLAN to **200**
- Set the MEP direction to **down** (optional because down is the direction default)
- Enable the sending of CCM messages for MEP **203**
- Set the remote MEP that CCM messages will be sent to **201** and **202**
- Activate the MEP state machines
- Enable the MEP configuration

### Device maCE2:3 CLI Input

```

maCE2:3 (rw-config) ->cfm enable
maCE2:3 (rw-config) ->cfm logging filter md string-name mdCE1 ma string-name maCE2
maCE2:3 (rw-config) ->cfm md string-name mdCE1
maCE2:3 (rw-config-cfm-md.1) ->level 5
maCE2:3 (rw-config-cfm-md.1) ->enable
maCE2:3 (rw-config-cfm-md.1) ->ma string-name maCE2
maCE2:3 (rw-config-cfm-ma.1) ->mep-list 201-203 enable
maCE2:3 (rw-config-cfm-ma.1) ->enable
maCE2:3 (rw-config-cfm-ma.1) ->ma-comp
maCE2:3 (rw-config-cfm-macomp) ->vid 200
maCE2:3 (rw-config-cfm-macomp) ->enable
maCE2:3 (rw-config-cfm-macomp) ->exit
maCE2:3 (rw-config-cfm-ma.1) ->mep 201
maCE2:3 (rw-config-cfm-mep.102) ->port tg.1.6
maCE2:3 (rw-config-cfm-mep.102) ->vid 200
maCE2:3 (rw-config-cfm-mep.102) ->direction down
maCE2:3 (rw-config-cfm-mep.102) ->cci-enabled
maCE2:3 (rw-config-cfm-mep.102) ->remote-mep active 201,203
maCE2:3 (rw-config-cfm-mep.102) ->active
maCE2:3 (rw-config-cfm-mep.102) ->enable
maCE2:3 (rw-config-cfm-mep.102) ->exit

```

```

maCE2:3 (rw-config-cfm-ma.1) ->exit
maCE2:3 (rw-config-cfm-ma.1) ->exit
maCE2:3 (rw-config) ->

```

## Configuring Switch 2

To configure CE Device 2:

- Enable CFM in global configuration mode
- Limit logging to MA **maCE2**
- Configure the default MD VLAN to **100**
- Configure the default MD VLAN 100 MD level to **5**
- Set the default MD VLAN 100 MHF create setting to **default**

### Switch 1 CLI Input

```

Switch 2 (rw-config) -> cfm enable
Switch 2 (rw-config) -> cfm logging filter md string-name mdCE2 ma string-name maCE1
Switch 2 (rw-config) -> cfm default-md vid 100
Switch 2 (rw-config-def-md.100) -> level 5
Switch 2 (rw-config-def-md.100) -> mhf-creation default
Switch 2 (rw-config-def-md.100) -> exit
Switch 2 (rw-config) ->

```

## Multiple MD Configuration Example

A CFM configuration requires multiple MDs when one or more network segments of a monitored service are either not directly under the control of a single administrator or the administrator chooses to segment the network for administrative purposes. Our multiple MD configuration example is a basic customer equipment VLAN monitored service that passes through the service provider MD and two network operator MDs. See the [“Single MD Configuration Example”](#) on page 27-31 if you are using CFM in a single administrative control network environment.

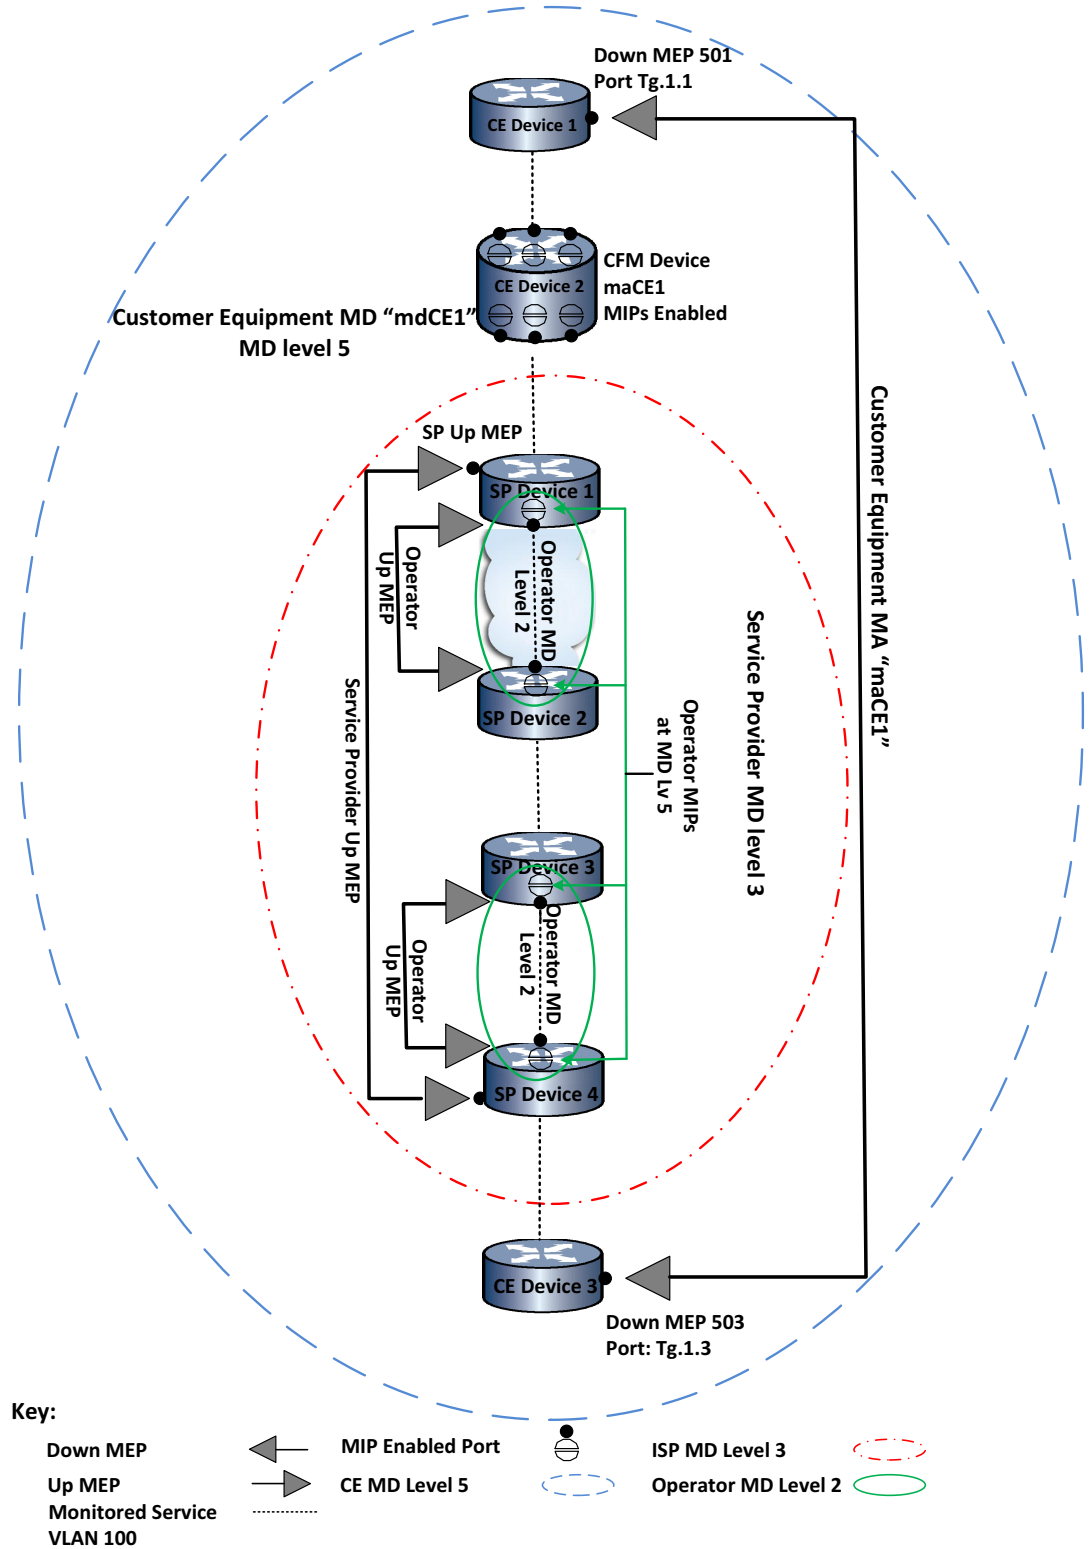
The example will provide the customer equipment configuration for the two monitored service end points (CE Device 1 and CE device 3) and a MIP enabled device (CE Device 2) that will snoop CCMs that traverse the service between the two end points. CE Device 2 configuration is representative of all devices within the customer equipment MD (level 5) that reside between the two end points with the exception that enabling MIPs is optional.

See [Figure 27-7](#) on page 27-40 for a presentation of the multiple MD configuration. The configuration of the service provider and network operator MDs is transparent to the customer equipment MD and is discussed only to the extent that it relates to understanding the customer equipment configuration.

In our example, CFM monitors VLAN 100. The customer equipment resides within MD level 5. The service provider equipment resides within MD 3, and the two network operators' equipment resides within two MD level 2 domains. The down MEPs 501 and 503 are the two customer service end points for the service. MIPs are enabled on customer equipment CE Device 2 to snoop CCMs that traverse the device. The customer equipment administrator has requested that the network operators enable MIPs on their Up MEPs configured for MD level 5 providing snooping of CCMs from MEPs 501 and 503 that transit the service. This configuration is performed by the network operator administrators. MIPs are not enabled on service provider ports to prevent exposing the customer service network topology within the service provider network.



Figure 27-7 Multiple MD Example Configuration Overview



## Configuring CE Device 1

To configure CE Device 1:

- Enable CFM in global configuration mode
- Limit logging to MA **maCE1**
- Access MD mode for MD **mdCE1** using the string-name naming convention
- Set the MD level to **5**
- Enable the MD configuration
- Access MA mode for MA **maCE1** using the string-name naming convention
- Set the interval between sending CCM messages to **10** seconds
- Configure the MEP list for the two down MEPs: **501** and **503**
- Enable the MA configuration
- Access the MA-Comp mode for **maCE1**
- Set VLAN 100 as the monitored service
- Enable MA-Comp configuration for **maCE1**
- Access MEP configuration for MEP **501**
- Set the MEP port to **tg.1.1**
- Set the MEP VLAN to **100**
- Set the MEP direction to **down** (optional because down is the direction default)
- Enable the sending of CCM messages for MEP **501**
- Set the remote MEP that CCM messages will be sent to **503**
- Activate the MEP state machines
- Enable the MEP configuration

## CE Device 1 CLI Input

```

CE Device 1(rw-config)->cfm enable
CE Device 1(rw-config)->cfm logging filter md string-name mdCE1 ma string-name
maCE1
CE Device 1(rw-config)->cfm md string-name mdCE1
CE Device 1(rw-config-cfm-md.1)->level 5
CE Device 1(rw-config-cfm-md.1)->enable
CE Device 1(rw-config-cfm-md.1)->ma string-name maCE1
CE Device 1(rw-config-cfm-ma.1)->ccm-interval 10sec
CE Device 1(rw-config-cfm-ma.1)->mep-list 501,503 enable
CE Device 1(rw-config-cfm-ma.1)->enable
CE Device 1(rw-config-cfm-ma.1)->ma-comp
CE Device 1(rw-config-cfm-macomp)->vid 100
CE Device 1(rw-config-cfm-macomp)->enable
CE Device 1(rw-config-cfm-macomp)->exit
CE Device 1(rw-config-cfm-ma.1)->mep 501
CE Device 1(rw-config-cfm-mep.501)->port tg.1.1
CE Device 1(rw-config-cfm-mep.501)->vid 100
CE Device 1(rw-config-cfm-mep.501)->direction down
CE Device 1(rw-config-cfm-mep.501)->cci-enabled

```

```
CE Device 1(rw-config-cfm-mep.501)->remote-mep active 503
CE Device 1(rw-config-cfm-mep.501)->active
CE Device 1(rw-config-cfm-mep.501)->enable
CE Device 1(rw-config-cfm-mep.501)->exit
CE Device 1(rw-config-cfm-ma.1)->exit
CE Device 1(rw-config-cfm-ma.1)->exit
CE Device 1(rw-config)->
```

## Configuring CE Device 2

To configure CE Device 2:

- Enable CFM in global configuration mode
- Limit logging to MA **maCE1**
- Access MD mode for MD **mdCE1** using the string-name naming convention
- Set the MD level to **5**
- Enable the MD configuration
- Access MA mode for MA **maCE1** using the string-name naming convention
- Enable the MA configuration
- Access the MA-Comp mode for **maCE1**
- Set VLAN 100 as the monitored service
- Enable MIP creation for all ports on the device
- Enable MA-Comp configuration for **maCE1**

### CE Device 2 CLI Input

```
CE Device 2(rw-config)->cfm enable
CE Device 2(rw-config)->cfm logging filter md string-name mdCE1 ma string-name
maCE1
CE Device 2(rw-config)->cfm md string-name mdCE1
CE Device 2(rw-config-cfm-md.1)->level 5
CE Device 2(rw-config-cfm-md.1)->enable
CE Device 2(rw-config-cfm-md.1)->ma string-name maCE1
CE Device 2(rw-config-cfm-ma.1)->enable
CE Device 2(rw-config-cfm-ma.1)->mep-list 501,503 enable
CE Device 2(rw-config-cfm-ma.1)->ma-comp
CE Device 2(rw-config-cfm-macomp)->vid 100
CE Device 2(rw-config-cfm-macomp)->mhf-creation default
CE Device 2(rw-config-cfm-macomp)->enable
CE Device 2(rw-config-cfm-macomp)->exit
CE Device 2(rw-config-cfm-ma.1)->exit
CE Device 2(rw-config-cfm-md.1)->exit
CE Device 2(rw-config)->
```

## Configuring CE Device 3

To configure CE Device 3:

- Enable CFM in global configuration mode
- Limit logging to MA **maCE1**
- Access MD mode for MD **mdCE1** using the string-name naming convention
- Set the MD level to **5**
- Enable the MD configuration
- Access MA mode for MA **maCE1** using the string-name naming convention
- Set the interval between sending CCM messages to **10** seconds
- Configure the MEP list for the two down MEPs: **501** and **503**
- Enable the MA configuration
- Access the MA-Comp mode for **maCE1**
- Set VLAN 100 as the monitored service
- Enable MA-Comp configuration for **maCE1**
- Access MEP configuration for MEP **503**
- Set the MEP port to **tg.1.3**
- Set the MEP VLAN to **100**
- Set the MEP direction to **down** (optional because down is the direction default)
- Enable the sending of CCM messages for MEP **503**
- Set the remote MEP that CCM messages will be sent to **501**
- Activate the MEP state machines
- Enable the MEP configuration

### CE Device 3 CLI Input

```
CE Device 3(rw-config)->cfm enable
CE Device 3(rw-config)->cfm logging filter md string-name mdCE1 ma string-name
maCE1
CE Device 3(rw-config)->cfm md string-name mdCE1
CE Device 3(rw-config-cfm-md.1)->level 5
CE Device 3(rw-config-cfm-md.1)->enable
CE Device 3(rw-config-cfm-md.1)->ma string-name maCE1
CE Device 3(rw-config-cfm-ma.1)->ccm-interval 10sec
CE Device 3(rw-config-cfm-ma.1)->mep-list 501,503 enable
CE Device 3(rw-config-cfm-ma.1)->enable
CE Device 3(rw-config-cfm-ma.1)->ma-comp
CE Device 3(rw-config-cfm-macomp)->vid 100
CE Device 3(rw-config-cfm-macomp)->enable
CE Device 3(rw-config-cfm-macomp)->exit
CE Device 3(rw-config-cfm-ma.1)->mep 503
CE Device 3(rw-config-cfm-mep.501)->port tg.1.3
```

```

CE Device 3(rw-config-cfm-mep.501)->vid 100
CE Device 3(rw-config-cfm-mep.501)->direction down
CE Device 3(rw-config-cfm-mep.501)->cci-enabled
CE Device 3(rw-config-cfm-mep.501)->active
CE Device 3(rw-config-cfm-mep.501)->remote-mep active 501
CE Device 3(rw-config-cfm-mep.501)->enable
CE Device 3(rw-config-cfm-ma.1)->exit
CE Device 3(rw-config-cfm-ma.1)->exit
CE Device 3(rw-config)->

```

## Terms and Definitions

Table 27-9 lists terms and definitions used in this CFM configuration discussion.

**Table 27-9 Connectivity Fault Management (CFM) Terms and Definitions**

Term	Definition
CFM Linktrace Protocol	A diagnostic protocol used to verify the path between the initiating device and a target device, and to help determine where in the path a problem might exist.
CFM Loopback Protocol	A diagnostic protocol used to determine connectivity between the initiating device and the target device.
CFM monitored service	The network entity monitored by CFM (currently a single VLAN per MA).
Connectivity Check Messages (CCM)	A CCM is a unidirectional multicast message, confined to a single operator domain that provides a means to detect connectivity failures in an MA. These messages are unidirectional and do not solicit a response. Each MEP transmits a periodic multicast CCM inward towards the other MEPs belonging to the operator.
Connectivity Fault Management (CFM)	Provides network operators the means to monitor and troubleshoot services that may span multiple domain Ethernet networks based upon IEEE 802.1Q-2011.
ID Permission	Specifies whether chassis or remote management information or both are sent in the SenderID TLV.
Maintenance Association (MA)	A logical container for a specific CFM monitored service.
Maintenance Domain (MD)	A logical container for all the equipment associated with the CFM monitored service and owned by a single network operator.
Maintenance End-Point (MEP)	A port, belonging to an MA, through which data enters and exits the portion of the network monitored by the CFM service.
Maintenance Intermediate-Point (MIP)	An auto-created MP on a port that resides along the path between MEPs, supplements the MEP functionality by passively snooping the CCMs that pass through it, and is able to respond to a loopback or linktrace message.
Maintenance Point (MP)	A demarcation point on a port that implements the CFM functions within a MA.
MD level	A CFM MD value that determines the reach and scope of the organization controlling the MD, with higher levels able to encapsulate lower levels, allowing CCMs to pass transparently through lower level MDs.
MEP direction	Specifies whether a MEP faces the link or the bridge relay thereby controlling the transmission direction of CFM PDUs by the MEP.
MEP list	A list containing all the MEPs in an MA.

**Table 27-9 Connectivity Fault Management (CFM) Terms and Definitions (continued)**

<b>Term</b>	<b>Definition</b>
MHF creation	Specifies whether MIPs are created on the ports for the configured device.
VLAN table	The association of one or more CFM services with a primary CFM service.



## Virtual Routing and Forwarding (VRF) Configuration

This document provides the following information about configuring Virtual Routing and Forwarding (VRF) on the Extreme Networks 7100-Series platforms.

For information about...	Refer to page...
<a href="#">Using VRF in Your Network</a>	28-1
<a href="#">Implementing VRF</a>	28-1
<a href="#">VRF Overview</a>	28-1
<a href="#">Configuring VRF</a>	28-4
<a href="#">Terms and Definitions</a>	28-4

### Using VRF in Your Network

Virtual Routing and Forwarding (VRF) provides a method of partitioning your network into different routed domains. A VRF is a segregated domain for the routed forwarding of packets. VRFs are used to divide a router into multiple standalone forwarding domains that may contain unique IP networks, routes, and other configuration that would otherwise conflict if they were all deployed on the same router. VRFs can exchange routes between one another. An Interface may be configured to one and only one VRF. An interface configured to a particular VRF is considered a member of that VRF. One or more VRF(s) can be used as a gateway (or access point) to a larger Internet.

### Implementing VRF

To configure a VRF:

- Create the VRF in any command mode, optionally specifying an SNMPv3 context name.
- Enter the VRF router mode, followed by entering configuration mode for the created VRF.

### VRF Overview

For information about...	Refer to page...
<a href="#">VRFs, Interfaces, and IP Addresses</a>	28-3



7100-Series devices have a single default router named “global”. The global router:

- Exists when you first boot the device
- Manages the VRFs for this physical router
- Can neither be created nor deleted
- Can manage up to 16 VRF instances depending upon your system

Use the **show limits application vrf** command to determine the number of VRF instances your system supports.

Each optional VRF instance you create functions as its own routing domain. All routing features and protocols that are supported on the global router are also supported in a VRF instance. VRF instance router protocol configuration (for example, configuring PIM, OSPF, and IGMP) is identical to the global router protocol configuration. Protocol configurations in different VRFs do not conflict with each other because they are completely separate instances of the protocol.

You create a VRF router, in any command mode, using the **set router vrf create** command. The command requires that you specify a name of up to 31 printable characters, except for the space character. Extreme Networks recommends that you provide the VRF with a meaningful name such as “Marketing” or “Internet-Access”.

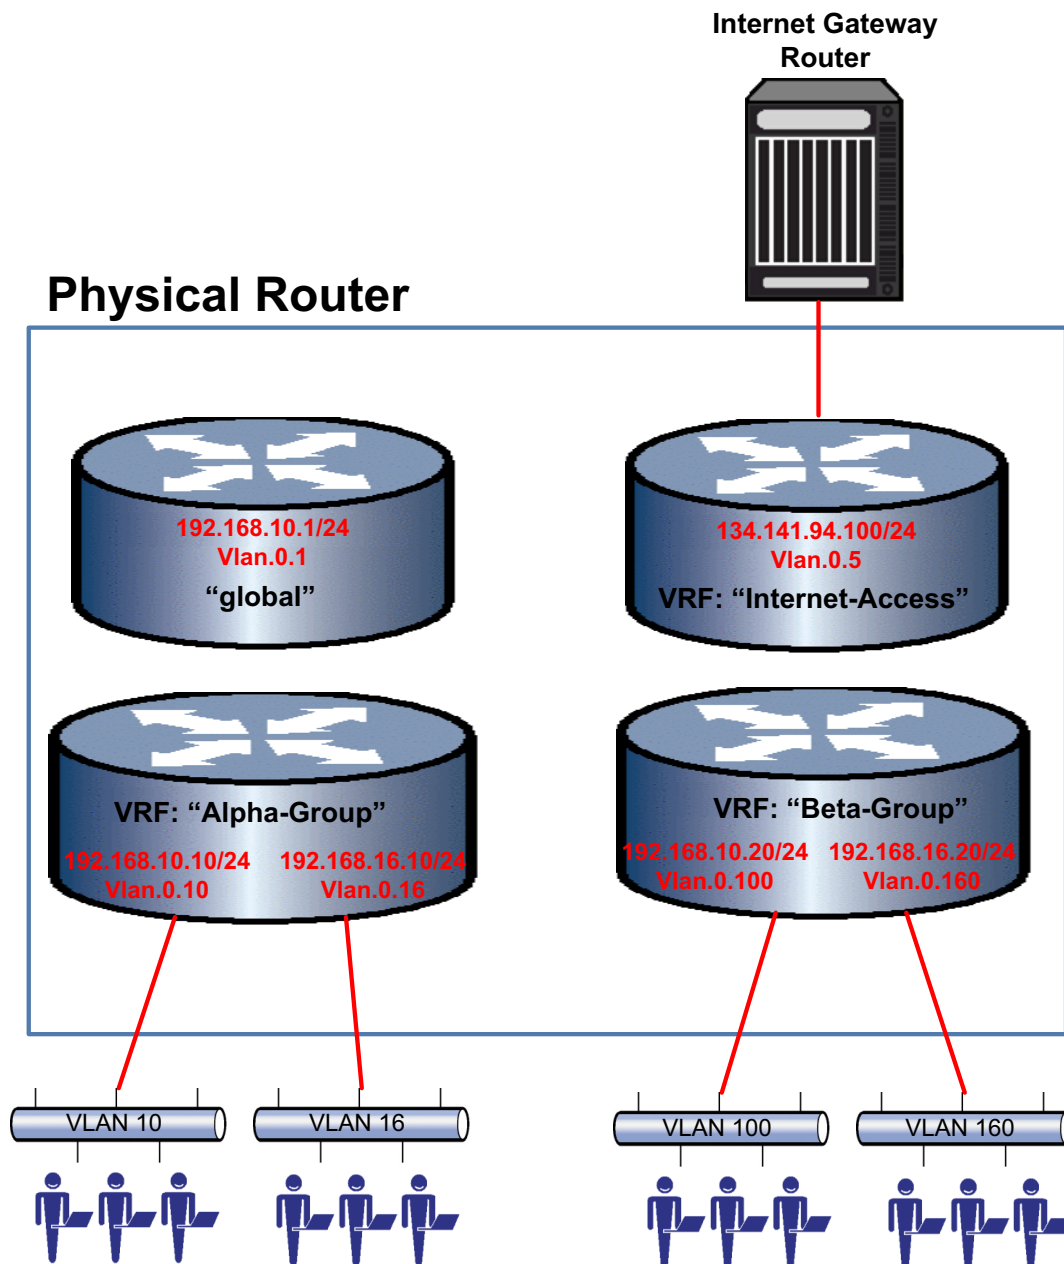
You can optionally specify an SNMPv3 context of up to 28 characters. If not specified, the SNMPv3 context defaults to the name of the VRF instance. If the VRF instance name exceeds 28 characters, the SNMPv3 context must be specified when creating the VRF. Refer to the **set router vrf create** command for information on creating a VRF instance.

The behavior when clearing the global router is different versus clearing a VRF instance. When you clear the global router, a blank configuration file is written to persistent memory. The global router is not deleted. Unlike the global router, all VRFs can be both created and deleted. When you clear a VRF, the VRF is deleted along with all of its configuration.

Use the **clear router vrf** command to clear the global router configuration or to delete a VRF instance from the system.

[Figure 28-1](#) on page 28-3 presents a router that has been segmented into three VRF routers: two VRF routers with user group access named Alpha-Group and Beta-Group, and a VRF for internet access named Internet-Access.

Figure 28-1 VRF Overview



## VRFs, Interfaces, and IP Addresses

By default, interfaces do not belong to any VRF instance until they are assigned. An interface may belong to only one VRF at a time. When you first create a VRF, the next available loopback interface is assigned as the default interface for the VRF router. Once bound to a VRF router, interfaces are configured in that VRF router context. You must first remove the bound VRF interface from its current VRF instance before moving the interface to a different VRF instance. To remove an interface from a VRF instance, along with all its configuration, use the command **no interface interface-name**.

In VRF configuration mode, the **interface interface-name** command automatically binds the named interface to the current VRF and enters interface configuration mode. If the interface has already been bound to a different VRF, an error message is displayed.

IP addresses assigned in different VRFs are completely separate, thus overlapping or identical IP addressing is permitted across different VRFs. For example, VRF “Corporate” may have IP address range 10.1.100.1/16 associated with interface ge.1.1 while the “Marketing” VRF has IP address range 10.1.100.1/16 associated with interface ge.1.2. As a packet ingresses the router, the interface it ingresses on will determine which VRF router will receive it.

## Configuring VRF

This section provides details for the configuration of VRF on the 7100-Series products.

[Table 28-1](#) lists VRF parameter default values.

**Table 28-1 Default VRF Parameters**

Parameter	Description	Default Value
SNMPv3 context Name	The name that SNMPv3 will associate with this VRF.	VRF Name
router context	The VRF router command mode context if no router is specified	global

[Procedure 28-1](#) describes how to configure VRF.

**Procedure 28-1 VRF Configuration**

Step	Task	Command(s)
1.	Create the VRF, in any configuration mode, optionally specifying an SNMPv3 context name.	<b>set router vrf create</b> <i>vrf-name</i> [ <b>context</b> <i>context-name</i> ]
2.	Enter router mode for the VRF to be configured.	<b>router</b> [ <i>name</i> ]
3.	Enter configuration mode for this VRF router instance.	<b>configure</b>
4.	Optionally, when creating a policy route map, with a match IP address policy in which the interface belongs to a different VRF, configure the next hop VRF to perform the route lookup using its routing table.	<b>set vrf</b> <i>vrf-name</i>

## Terms and Definitions

[Table 28-2](#) lists terms and definitions used in this VRF configuration discussion.

**Table 28-2 VRF Configuration Terms and Definitions**

Term	Definition
global router	The default router for the physical router. Also responsible for managing VRFs configured on the physical router.
SNMPv3 context	Specifies the SNMPv3 context name to be used by SNMP for a given VRF instance.
Virtual Routing and Forwarding (VRF)	A method of partitioning a global router network into different routed domains.
VRF instance	A segregated routing domain for the routed forwarding of packets managed by the global router.





## IP Routing Configuration

This document describes IPv4 and IPv6 routing configuration on Extreme Networks 7100-Series devices.

For information about...	Refer to page...
<a href="#">The Router</a>	<a href="#">29-1</a>
<a href="#">The Routing Interface</a>	<a href="#">29-3</a>
<a href="#">IP Static Routes</a>	<a href="#">29-12</a>
<a href="#">IPv6 Neighbor Discovery</a>	<a href="#">29-13</a>
<a href="#">Configuring IPv6 Neighbor Discovery</a>	<a href="#">29-14</a>
<a href="#">The ARP Table</a>	<a href="#">29-15</a>
<a href="#">Router Management and Information Display</a>	<a href="#">29-18</a>
<a href="#">Terms and Definitions</a>	<a href="#">29-20</a>

### The Router

The current firmware implementation supports a single default Virtual Routing and Forwarding (VRF) router named **global**.

There are two ways of accessing the **global** VRF router configuration:

- Directly from global configuration mode, accessed by entering the **configure** command from the system command mode
- First entering router command mode from system command mode using the **router** command, specifying **global** as the name of the router, followed by entering the **configure** command to gain access to the router configuration mode

Once in either router configuration or global configuration command mode, the same set of router configuration commands are available to you.

### Entering Router Configuration

To enter the **global** VRF router configuration context from the system command mode, and verify the current router context, enter:

```
7100 Chassis(rw)->configure
7100 Chassis(rw--config)->show router
Router Services are currently running on module 1.
VRF Context      : global
```

```
RD                : not set
7100 Chassis(rw-router-config)->

To enter the global VRF router configuration context from the router configuration command mode, and verify the current router context, enter:

7100 Chassis(rw)->router global
7100 Chassis(rw-router)->configure
7100 Chassis(rw-router-config)->show router
Router Services are currently running on module 1.
VRF Context      : global
RD                : not set
7100 Chassis(rw-router-config)->
```

[Table 29-1](#) describes how to enter router configuration mode.

**Table 29-1 Entering Router Configuration Mode**

Task	Command(s)
In system configuration mode, enter router command mode for the specified router.	<b>router</b> [name]
Supported routers: <b>global</b> .	
To enter router configuration command mode for the <b>global</b> or named VRF router, use the <b>configure</b> command in router command mode.	
The <b>global</b> router can also be configured in global configuration command mode.	

## Display Router Configuration

Use the **show router** command in any command mode to display router settings for the current VRF context.

Use the **show limits** command in any command mode to display application limits associated with the current VRF context. Use the **show limits application** command to display the limits for a specified application in the current VRF context. Use the **show limits resource-ipv6netmask** command to display the IPv6 netmask setting. The following example displays a sample output of the **show limits** command:

```
7100 Chassis(su)->show limits

Chassis limits:
Application                               Limit    In use   Entry size  Total Memory
-----
ip-interfaces                             1108     3        -           -
  lo-interfaces                             1        1        -           -
  lpbk-interfaces                           33        0        -           -
  vlan-interfaces                          1024     2        -           -
static-routes                             2048     1        176B       352K
Total Memory                               -        -        -           352K
7100 Chassis(su)->
```

Use the **show running-config** command to display non-default router configuration for either all or a specified option. When specifying **all**, both default and non-default configuration displays.

Additional options are available for the display of a subset of the running configuration by feature or protocol. Enter the **show running-config ?** command for a listing of the additional options. The following example displays a sample output of the **show running-config** command:

```
7100 Chassis(su)->show running-config
# **** Global Router Configuration ****
configure terminal
!
interface vlan.0.1
 ip address 100.10.10.10 255.0.0.0 primary
 ip dvmrp
 no shutdown
 exit
interface vlan.0.56
.
.
.
7100 Chassis(su)->
```

## The Routing Interface

For information about...	Refer to page...
<a href="#">IP Routing Addresses</a>	<a href="#">29-4</a>
<a href="#">Secondary and Private VLAN</a>	<a href="#">29-7</a>
<a href="#">Non-Forwarding IP Management Interfaces</a>	<a href="#">29-9</a>
<a href="#">Show Interface Examples</a>	<a href="#">29-10</a>

Routing interfaces are configured by entering the **interface** command from the configuration command mode, specifying the interface ID and whether the interface is a VLAN or a loopback interface. If the interface has not previously been created, this command creates a new routing interface.

A VLAN routing interface can be configured before its VLAN is created in system configuration mode, but VLANs must be created from the system CLI before they will be operational within IP routing. See “[Configuring VLANs](#)” on page 24-9 for VLAN configuration details.

Each VLAN or loopback interface must be configured for routing separately using the **interface** command. To end configuration on one interface before configuring another, type **exit** at the command prompt. Enabling the interface for IP routing is required using the **no shutdown** command before exiting the interface mode.

IPv4 forwarding is enabled by default on a routing interface. Use the **no ip forwarding** command within interface configuration command mode to disable IPv4 forwarding on a routing interface.

IPv6 forwarding is disabled by default on a routing interface.



**Note:** IPv4 and IPv6 forwarding are both enabled by default on loopback interfaces. Without forwarding, a loopback interface is unreachable. This configuration setting cannot be modified.



## IP Routing Addresses

### IPv4 Interface Address

A single primary network IPv4 address is configurable on an interface. Up to 100 secondary network IPv4 addresses are configurable. The first network IP address assigned to an interface is the primary whether explicitly configured as primary or not. To configure a secondary network IP address on an interface, the address must be explicitly configured as secondary, otherwise you will be queried as to whether you want to overwrite the current primary.

In the following example the IP address is set to **99.0.0.1/24**. This setting is followed by an attempt to configure **99.0.0.2/16** as a secondary address, while failing to specify the **secondary** keyword. When queried as to whether the primary IP address should be changed, **n** is entered. The **secondary** keyword is added on the next line. The **show running-config** command output confirms the configuration:

```
7100 Chassis(rw-config-intf-vlan.0.99)->ip address 99.0.0.1/24
7100 Chassis(rw-config-intf-vlan.0.99)->ip address 99.0.0.2/16
Do you want to replace primary IP address (y/n) [n]?n
7100 Chassis(rw-config-intf-vlan.0.99)->ip address 99.0.0.2/16 secondary
7100 Chassis(rw-config-intf-vlan.0.99)->show running-config interface vlan.0.99
# **** VRF default (default) ****
configure terminal
!
 interface vlan.0.99
   ip address 99.0.0.1 255.255.255.0 primary
   ip address 99.0.0.2 255.255.0.0 secondary
 exit
!
exit
7100 Chassis(rw-config-intf-vlan.0.99)->
```

The **ip address** command in interface configuration command mode is used to assign IP networks as primary or secondary to a routing interface.

See “[IPv6 Interface Address](#)” on page 29-5 for IPv6 address configuration information.

The **no ip address** command removes the specified IPv4 address configuration for this interface.

### IPv4 Router Interface Configuration Example

The following example:

- Creates the interface for VLAN 1
- Configures a primary IP address of 10.21.130.59 255.255.128.0

```
7100 Chassis(rw)->configure
7100 Chassis(rw-config)->interface vlan 1
7100 Chassis(rw-config-intf-vlan.0.1)->ip address 10.21.130.59 255.255.128.0
7100 Chassis(rw-config-intf-vlan.0.1)->no shutdown
7100 Chassis(rw-config-intf-vlan.0.1)->exit
7100 Chassis(rw-config)->
```

See the current firmware release notes for the number of routing interfaces supported on an 7100-Series routing module.

A primary IP address must be configured on each routing interface. Secondary IP addresses can optionally be configured. See the current firmware release notes for the number of secondary addresses supported on an interface and module. Use the **ip address** command in interface configuration command mode to assign an IP address and optional secondary IP addresses to an interface, specifying whether the assigned address is primary or secondary.

## IPv6 Interface Address

One or more unicast IPv6 addresses and a single link local address can be configured for an interface using the **ipv6 address** command in interface configuration mode.

Link local addresses are network addresses which are intended only for communications within one segment of a local network (a link) or a point-to-point connection. They allow addressing hosts without using a globally-routable address prefix. Routers will not forward packets with link-local addresses. A link local address must begin with **fe80:**.

An interface can be configured to have its IPv6 address auto acquired using the **autoconfig** option.

A single link local address is supported per interface. If IPv6 autoconfiguration is enabled, the link local address is autoconfigured. When manually configuring a link local address, if a link local address already exists on the interface, a warning displays asking you if you wish to change it.

EUI-64 is an IPv6 address automatic interface addressing capability. By implementing the IEEE's 64-bit Extended Unique Identifier (EUI-64) format, a host can automatically assign itself a unique 64-bit IPv6 interface identifier without the need for manual configuration or DHCP. This is accomplished on Ethernet interfaces by referencing the already unique 48-bit MAC address and reformatting that value to match the EUI-64 specification as specified in RFC 2373. When configuring an EUI-64 address, the specified prefix must have a length of 64.

A general prefix allows an assigned name to represent a network prefix from which longer IPv6 addresses can be configured. The sub-bits added to the general prefix can both extend the network prefix by adding to the specified prefix length, as well as complete the IPv6 address.

Use the **ipv6 general-prefix** command to configure a general prefix. See [“IPv6 General Prefix”](#) on page 29-5 for general prefix details.

Use the `show ipv6 interface` command to display IPv6 addresses assigned by the **ipv6 address** command.

See [“IPv4 Interface Address”](#) on page 29-4 for IPv4 address configuration information.

The **no ipv6 address** command removes the specified IPv6 address configuration for this interface.

## IPv6 General Prefix

The general prefix is an ease of use feature that allows an assigned name to represent a network prefix from which longer IPv6 addresses can be configured. Network renumbering is simplified by redefining the general prefix, thereby changing the portion of addresses to which the general prefix is assigned.

When using a general prefix to configure an IPv6 address, you can extend the network prefix by adding to the length specified in the **ipv6 address** command.

Deleting a general prefix does not delete the underlying addresses defined by the general prefix. Any IPv6 addresses based upon the general prefix remain. Use the **no ipv6 address** command to remove the IPv6 address.

The 7100-Series supports the configuration of up to 64 general prefixes on a system.

The following example creates a general prefix named “Doc-Prefix” with a prefix value of **2001:11ac:fd34::/48** and assigns the IPv6 address **2001:11ac:fd34:50:0:0:abcd:33** to VLAN 51. The general prefix **Doc-Prefix** is followed by **::50:0:0:abcd:33/64**. The subnet length is changed to **/64** adding **:50** to the general prefix to create a network prefix of **2001:11ac:fd34:50/64** for this IPv6 address:

```
7100 Chassis(su)->configure
7100 Chassis(su-config)->ipv6 general-prefix Doc-Prefix 2001:11ac:fd34::/48
7100 Chassis(su-config)->show ipv6 general-prefix
    ipv6 general-prefix Doc-Prefix 2001:11ac:fd34::/48
7100 Chassis(su-config)->interface vlan 51
7100 Chassis(su-config-intf-vlan.0.51)->ipv6 address Doc-Prefix
::50:0:0:abcd:33/64
7100 Chassis(su-config-intf-vlan.0.51)->show ipv6 interface vlan.0.51

vlan.0.51 is Operationally down, Administratively down
IPv6 is enabled link-local address is fe80::211:88ff:fe7c:32c1%vlan.0.51
Global unicast address(es):
    2001:11ac:fd34:50:0:0:abcd:33, subnet is 2001:11ac:fd34:50::/64
...
7100 Chassis(su-config-intf-vlan.0.51)->
```

## IPv6 Router Interface Configuration Examples

This example sets the IPv6 address for interface VLAN 50 to **ba10:1100:aa11:c171:0:0:1111:00/48**:

```
7100 Chassis(su-config)->interface vlan 50
7100 Chassis(su-config-intf-vlan.0.50)->ipv6 address
ba10:1100:aa11:c171:0:0:1111:00/48
7100 Chassis(su-config-intf-vlan.0.50)->
```

This example sets the IPv6 link local address for interface VLAN 50 to **fe80:1234:5678::300**:

```
7100 Chassis(su-config)->interface vlan 50
7100 Chassis(su-config-intf-vlan.0.50)->ipv6 address fe80:1234:5678::300
link-local
Do you want to replace IPv6 link-local address (y/n) [n]?y
7100 Chassis(su-config-intf-vlan.0.50)->
```

This example sets an IPv6 EUI-64 address for interface VLAN 50 based upon the prefix **2001:febd:1234:0/64**, and displays the EUI-64 address in the interface output:

```
7100 Chassis(su-config)->interface vlan 50
7100 Chassis(su-config-intf-vlan.0.50)->ipv6 address 2001:febd:1234:0/64 eui-64
7100 Chassis(su-config-intf-vlan.0.50)->show ipv6 interface vlan.0.50

vlan.0.50 is Operationally down, Administratively down
IPv6 is enabled link-local address is fe80::2e0:63ff:fe6b:1d26%vlan.0.50
Global unicast address(es):
    2001:febd:1234::2e0:63ff:fe6b:1d26, subnet is 2001:febd:1234::/64 [EUI]
...
7100 Chassis(su-config-intf-vlan.0.50)->
```

## Secondary and Private VLAN

The secondary VLAN configuration on an IP Interface provides the ability to associate multiple L2 VLANs with one L3 IP interface. The secondary VLAN feature provides for the configuration of private VLANs by configuring ports on the secondary VLAN as private members. A secondary VLAN port can be configured as a private VLAN member by restricting the ports they can egress to using the **set vlan egress** command. Members of the private VLAN are connected hosts that share the IP interface of the primary VLAN, while at the same time are restricted from directly communicating with each other. Hosts on the primary VLAN, also referred to as the community VLAN, can communicate directly with hosts on both the primary and private VLANs.

When configuring members of the private VLAN, set both the secondary and primary VLAN constraint to shared, using the same constraint set ID. This setting assures that both the primary and secondary VLAN use the same FID. VLAN constraint is set using the **set vlan constraint** command.

Set ports on the primary VLAN as members of the egress list for all ports on both the primary and secondary VLANs. Set private member ports on the secondary VLAN as members of the egress list for all members of the primary VLAN. Use the **set vlan egress** command to set ports as members of a VLAN's egress list.

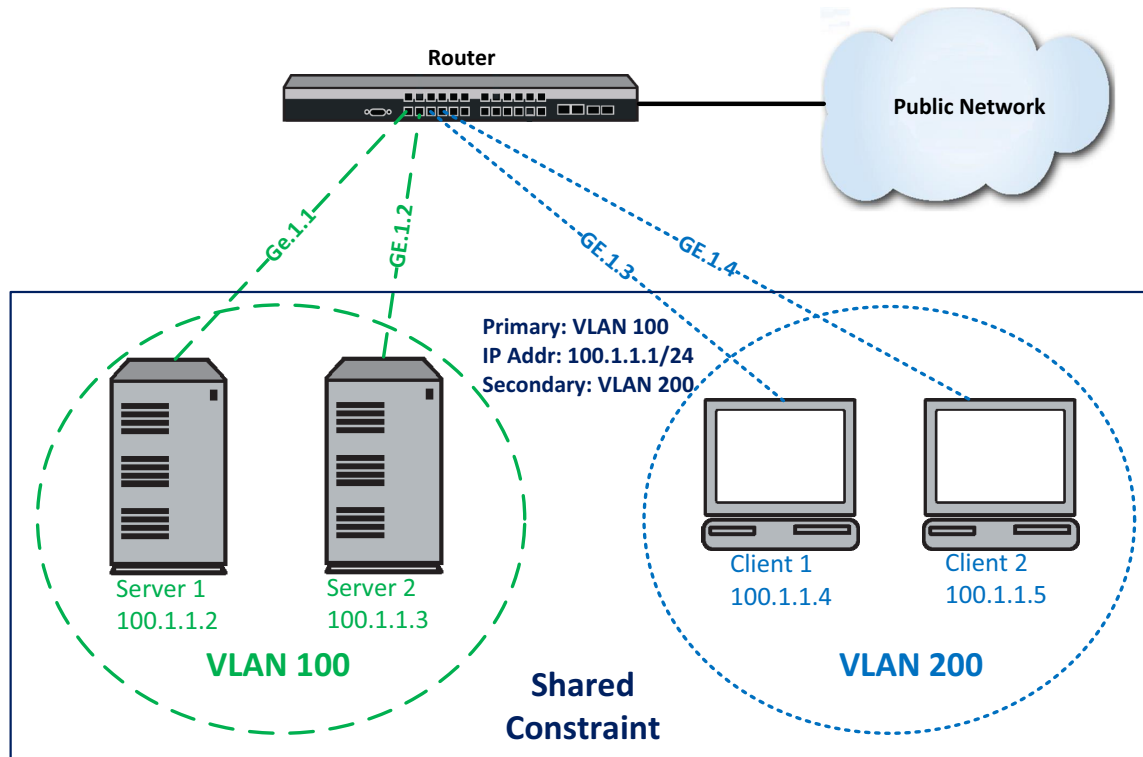
The secondary VLAN is not configured as an independent routing interface; it is configured within the primary VLAN. Only set an IP address for the primary VLAN interface. Do not set an IP address for the secondary VLAN.

Refer to [Chapter 24, VLAN Configuration](#) for VLAN configuration details.

This feature could be used by an internet service provider network where clients should not be directly communicating with other clients on the same network unless permitted to do so. These restricted clients would be assigned to the secondary VLAN.

### Private VLAN Configuration Example

[Figure 29-1](#) displays a private VLAN configuration example. VLAN 100 and VLAN 200 are VLANs configured on the 100.1.1.1/24 network. In this example VLAN 100 is the primary VLAN with members Server 1 and Server 2. VLAN 200 is the secondary VLAN with members Client 1 and Client 2. Primary VLAN members are configured on ports ge.1.1-2 and are members of the egress list for all ports on both VLAN 100 and VLAN 200. Private VLAN members are configured on secondary VLAN ports ge.1.3-4 with an egress only on VLAN 100 ports. Both the primary VLAN and the secondary VLAN are configured with the same constraint set ID of 100 which means they share the same filtering database (FID 100). The routing interface is VLAN 100. The secondary VLAN is configured within the routing interface VLAN 100 configuration mode.

**Figure 29-1 Secondary VLAN Configuration**

To configure this example:

- Create the static primary (VLAN 100) and secondary (VLAN 200) VLANs
- Assign ports **ge.1.1-2** to the primary VLAN
- Assign ports **ge.1.3-4** to the secondary VLAN
- Configure VLAN 200 as a private VLAN by:
  - Setting egress for VLAN 100 for all ports
  - Setting egress for VLAN 200 only on primary VLAN ports ge.1.1-2
- Set the VLAN constraint to shared for each VLAN with a constraint set ID of 100
- Configure the primary interface with a primary IP address of **100.1.1.1/24** and a secondary VLAN of **200**

```

7100 Chassis(rw)->set vlan name 100 PrimaryVlan
7100 Chassis(rw)->set vlan name 200 SecondaryVlan
7100 Chassis(rw)->set port vlan ge.1.1-2 100
7100 Chassis(rw)->set port vlan ge.1.3-4 200
7100 Chassis(rw)->set vlan egress 100 ge.1.1-4 untagged
7100 Chassis(rw)->set vlan egress 200 ge.1.1-2 untagged
7100 Chassis(rw)->set vlan constraint 100 100 shared
7100 Chassis(rw)->set vlan constraint 200 100 shared
7100 Chassis(rw)->configure
7100 Chassis(rw-config)->interface vlan 100
7100 Chassis(rw-config-intf-vlan.0.100)->ip address 100.1.1.1/24 primary
7100 Chassis(rw-config-intf-vlan.0.100)->secondary-vlan 200

```

## Non-Forwarding IP Management Interfaces

Multiple IP interface configuration provides the ability to assign a unique IP address to each non-routing interface on the switch. The ability to set a unique IP address on each VLAN configured on the switch means that host management can be accessed from any VLAN configured with its own IP address.

The ability to assign an IP subnet to an interface that is separate from the subnet that is passing data through the switch, allows the network administrator to create an out-of-band management subnet designed to only pass network management data.



**Note:** All interfaces can be configured as either a routing interface or a non-forwarding IP interface. It is recommended that you only use the non-routing multiple IP interface feature on a non-routing switch: a switch that does not have any routing capability turned on and is not directly connected to a router.

A non-routing host management IP interface can now be configured:

- In interface configuration command mode using the **interface** command
- In any command mode using an enhanced **set ip address** command

When configuring the non-routing host management IP interface in interface configuration command mode you must explicitly set the interface as a non-forwarding interface using the **no ip forwarding** command for IPv4 forwarding. IPv6 forwarding is disabled by default. On an IPv4 interface, you must disable IP Proxy ARP using the **no ip proxy-arp** command.

When configuring a non-routing host management IPv4 and IPv6 interfaces in any command mode, use the **set ip address** command. The IP address is assigned to the specified interface. The **set ip address** command automatically configures the specified interface to disable both IP forwarding and IP Proxy ARP for IPv4. IPv6 forwarding is disabled by default and IPv6 proxy is not supported. This command can only be used in a non-routing host management IP interface context.

The **set ip address** command only allows for the specifying of a primary IPv4 address or an IPv6 address. If you wish to configure a non-forwarding IP interface with secondary IP addresses, use the **interface** command in configuration command mode to configure the interface. IPv6 addressing makes no distinction between primary and secondary addresses and treats IPv6 addresses equally.

When clearing an IPv4 or IPv6 address, the IP address to be cleared is explicitly stated. This command can be used on a primary IPv4 address or any IPv6 address. Use the **no ip address** command in interface configuration command mode to clear a secondary IP address.

Use the **clear ip interface** command to clear the IP interface the IP address is assigned to.

The following example clears the IP interface VLAN 5:

```
7100 Chassis(rw)->clear ip interface vlan.0.5
```

### Non-Forwarding IPv4 Management Interface Examples

The following multiple IP interface example configures VLANs 1 and 5 as non-routing host management IP interfaces in interface configuration command mode. Both interfaces are configured with IP forwarding and IP Proxy ARP disabled as follows:

```
7100 Chassis(rw)->configure
7100 Chassis(rw-config)->interface vlan.0.1
7100 Chassis(rw-config-intf-vlan.0.1)->ip address 125.50.10.1/16
7100 Chassis(rw-config-intf-vlan.0.1)->no ip forwarding
7100 Chassis(rw-config-intf-vlan.0.1)->no ip proxy-arp
```

```
7100 Chassis(rw-config-intf-vlan.0.1)->no shutdown
7100 Chassis(rw-config-intf-vlan.0.1)->exit
7100 Chassis(rw-config)->interface vlan.0.5
7100 Chassis(rw-config-intf-vlan.0.5)->ip address 125.100.10.1/16
7100 Chassis(rw-config-intf-vlan.0.5)->no ip forwarding
7100 Chassis(rw-config-intf-vlan.0.5)->no ip proxy-arp
7100 Chassis(rw-config-intf-vlan.0.5)->no shutdown
7100 Chassis(rw-config-intf-vlan.0.5)->exit
7100 Chassis(rw-config)->
```

The above example is replicated below using the **set ip address** command in system command mode:

```
7100 Chassis(rw)->set ip address 125.50.10.1 mask 255.255.0.0 interface vlan.0.1
7100 Chassis(rw)->set ip address 125.100.10.1 mask 255.255.0.0 interface vlan.0.5
7100 Chassis(rw)->
```

## Non-Forwarding IPv6 Management Interface Examples

The following multiple IPv6 interface example configures VLANs 1 and 5 as non-routing host management IP interfaces in interface configuration command mode. IPv6 forwarding is disabled by default and IPv6 does not support proxy configuration:

```
7100 Chassis(rw-config)->interface vlan 1
7100 Chassis(rw-config-intf-vlan.0.1)->ipv6 address
ba10:1100:aa11:c171:0:0:1111:1/48
7100 Chassis(rw-config-intf-vlan.0.1)->no shutdown
7100 Chassis(rw-config-intf-vlan.0.1)->exit
7100 Chassis(rw-config)->interface vlan.0.5
7100 Chassis(rw-config-intf-vlan.0.5)->ipv6 address
ba10:1100:aa11:c171:0:0:111:5/48
7100 Chassis(rw-config-intf-vlan.0.5)->no shutdown
7100 Chassis(rw-config-intf-vlan.0.5)->exit
7100 Chassis(rw-config)->
```

## Setting a Default Host Management IP Interface

Setting the default host management interface is not supported in interface configuration command mode accessed using the **interface** command. In release 7.0, the **set ip interface** command can be entered in any command mode and provides for the optional setting of the interface as the default host management interface. The **set ip interface** command also allows for the initial configuration of a non-routing IP interface that you can assign an IP address to using the **set ip address** command.

## Show Interface Examples

Use the **show interface** command to display information about one or more VLAN or loopback interfaces configured on the router.

```
7100 Chassis(rw-config)->show interface vlan.0.1
vlan.0.1 is Operationally down, Administratively down
  MAC-Address is: 00-11-88-fe-79-68
  The name of this device is vlan.0.1
```

```
ARP/ND cache limit is 12288
MTU is 1500 bytes
Encapsulation ARPA, Loopback not set
IP Policy Routing not implemented
```

Use the **show ip interface** command to display information for interfaces configured for IP.

```
7100 Chassis(rw-config)->show ip interface vlan.0.1
vlan.0.1 is Operationally down, Administratively down
  IP forwarding disabled
  Frame Type ARPA
  Incoming IPv4 Access list is not set
  Outgoing IPv4 Access list is not set
  MTU is 1500 bytes
  ARP Timeout is 3600 seconds
  ARP Retransmit Time is 1 seconds
  ARP Stale-Entry-Timeout is 1200 seconds
  Proxy ARP is disabled
  Gratuitous ARP updating is set to update on ARP replies and ARP requests
  Gratuitous ARP learning is not set
  Multicast ARP learning is disabled
ICMP Mask Replies are always sent
  IP Policy Routing is not implemented
  PIM Sparse-mode is disabled
  DVMRP is disabled
```

This example shows how to display IPv6 configuration information for VLAN 51:

```
7100 Chassis(rw)->show ipv6 interface vlan.0.51

vlan.0.51 is Operationally down, Administratively down
  IPv6 is enabled link-local address is fe80::21f:45ff:fe5b:f5cf%vlan.0.51
  Global unicast address(es):
    2001:11ac:fd34:50::abcd:33, subnet is 2001:11ac:fd34:50::/64
  Joined group address(es):
    (None)
  IPV6 forwarding disabled
  IPV6 address auto-configuration is enabled
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ND DAD is enabled, number of DAD attempts: 1
7100 Chassis(rw)->
```



[Procedure 29-1](#) describes how to configure the routing interface.

### Procedure 29-1 Configuring the Routing Interface

Step	Task	Command(s)
1.	Enter router interface configuration command mode for the specified interface, from either global configuration or router configuration command mode.	<b>interface</b> {vlan <i>vlan-id</i>   loopback <i>loopback-id</i>   interface-name}
2.	Set the primary, and optionally the secondary or management, IPv4 address for this interface, in interface configuration command mode.	<b>ip address</b> { <i>ip-address</i>   <i>ip-address/prefixLength</i> } <i>ip-mask</i> [ <b>primary</b>   <b>secondary</b>   <b>management</b> ]
3.	Optionally, configure an IPv6 general prefix in global configuration mode to be assigned to an IPv6 address.	<b>ipv6 general-prefix</b> <i>name prefix/length</i>
4.	Set the IPv6 address for this interface in interface configuration command mode.	<b>ipv6 address</b> { <i>link-local-address link-local</i>   <i>ipv6-address/length</i>   <i>ipv6-prefix/length eui-64</i>   <b>autoconfig</b>   <i>general-prefix sub-bits/length</i> }
5.	Optionally disable IPv4 forwarding on this interface.	<b>no ip forwarding</b>
6.	Optionally, configure a secondary VLAN on the VLAN interface.	<b>secondary-vlan</b> <i>vlan-id</i>
7.	Enable this interface along with any changes made, in interface configuration command mode.	<b>no shutdown</b>

## IP Static Routes

An IP static route can be configured as a non-forwarding management route for IPv4.

Non-forwarding management routes can be configured using either the **ip route** or **ipv6 route** commands in configuration command mode or the **set ip route** in any command mode. See “[Traffic Non-Forwarding IP Static Routes](#)” on page 29-12 for a non-forwarding static route discussion.

### Traffic Non-Forwarding IP Static Routes

Non-forwarding IP static routes are management routes.

There are two methods for configuring a non-forwarding management route. The recommended method is to first set the routing interface as a non-forwarding interface using the IPv4 **no ip forwarding** command in interface configuration mode (IPv6 forwarding is disabled by default). In global configuration mode, configure the static route using the **ip route** command for an IPv4 route or **ipv6 route** command for an IPv6 route. Because the **ip route** and **ipv6 route** commands are configured in router configuration command mode, the configuration is capable of automatically determining the correct VLAN if not specified.

The second method is using the legacy command **set ip route** in system configuration mode specifying an IPv4 or IPv6 destination address.

### Traffic Non-Forwarding IP Static Route Examples

Non-forwarding interfaces are configured for IPv4 traffic using the **no ip forwarding** command and for IPv6 traffic using the **no ipv6 forwarding** command, in interface configuration mode. IPv6

forwarding is disabled by default on the interface. The following example enters static routes specifying the non-forwarding interface **vlan.0.4000** as the next-hop interface:

```
7100 Chassis(rw-config)->interface vlan.0.4000
7100 Chassis(rw-config-intf-vlan.0.4000)->no ip forwarding
7100 Chassis(rw-config-intf-vlan.0.4000)->exit
7100 Chassis(rw-config)->ip route 10.0.0.0/8 10.21.128.1 interface vlan.0.4000
7100 Chassis(rw-config)->ip route 125.20.0.0/16 125.20.10.1 interface vlan.0.4000
7100 Chassis(rw-config)->ipv6 route 2001:11ac:fd34::/48 2001:11ac:fd34:3333::4
interface vlan.0.4000
```

The following example uses the legacy method of configuring a non-forwarding static route from the system command mode with a destination of 192.122.173.42 and a gateway of 192.122.168.38:

```
7100 Chassis(rw)->set ip route 192.122.173.42 192.122.168.38
```

The following example uses the legacy method of configuring a non-forwarding static route from the system command mode with a destination of 192.122.173.50 and a next-hop interface of VLAN 50:

```
7100 Chassis(rw)->set ip route 192.122.173.50 vlan.0.50
```

[Procedure 29-2](#) describes how to configure a non-forwarding IP traffic route.

### Procedure 29-2 Configuring Non-forward IP Static Routes

Step	Task	Command(s)
1.	In interface configuration mode, set the routing interface for this static route to not forward IP traffic.	<b>no ip forwarding</b>
2.	In global configuration mode, configure the static route.	<b>ip route</b> { <i>prefix mask</i>   <i>prefix/prefix-length</i> }{ <i>ip-address</i> [ <b>recursive</b> ]   <b>interface</b> <i>interface-name</i>   <b>vlan</b> <i>vlan-id</i> } [ <i>distance</i> ] [ <b>tag</b> <i>tag-id</i> ] [ <b>blackhole</b> ] [ <b>reject</b> ]
3.	Optionally, in global configuration command mode, configure IPv6 static routes. IPv6 forwarding is disabled by default.	<b>ipv6 route</b> <i>prefix/length</i> { <i>ipv6-address</i> [ <b>recursive</b>   <b>interface</b> <i>interface-name</i> ]   <b>interface</b> <i>interface-name</i>   <b>vlan</b> <i>vlan-id</i>   <b>blackhole</b>   <b>reject</b> } [ <i>distance</i> ] [ <b>tag</b> <i>tag-id</i> ]
4.	Alternatively, in system configuration mode, configure the non-forwarding static route. This method supports legacy configurations. It is recommended that you use the method described in steps 1 - 3.	<b>set ip route</b> { <i>destination</i>   <b>default</b> } { <i>gateway</i>   <i>interface</i> } [ <i>mask</i> ]

## IPv6 Neighbor Discovery

The Neighbor Discovery (ND) protocol for IPv6 is defined in RFC4861. The neighbor discovery protocol uses ICMPv6 messages to determine the link-layer addresses of nodes residing on the same local link, to locate neighboring routers, to learn certain link and address configuration information, and to track the reachability of neighbors.

### Duplicate Address Detection

IPv6 Duplicate Address Detection (DAD) is described in RFC 4862. DAD uses neighbor solicitation and neighbor Advertisement messages to verify the uniqueness of an address. DAD

must be performed on unicast addresses prior to assigning them to an interface. An address remains in a tentative state while DAD is being performed. If a tentative address is found to be a duplicate, an error message is returned and the address is not assigned to the interface.

## IPv6 Address Autoconfiguration

IPv6 address autoconfiguration determines whether an interface IPv6 address will be auto-configured by acquiring the address from an attached router or must be manually configured. IPv6 address autoconfiguration is enabled for the interface by specifying the **autoconfig** option when entering the **ipv6 address** command in interface configuration mode.

The following example configures VLAN 3050 to acquire its IPv6 address from an attached router:

```
7100 Chassis(su-config)->interface vlan 3050
7100 Chassis(su-config-intf-vlan.0.3050)->ipv6 address autoconfig
7100 Chassis(su-config-intf-vlan.0.3050)->
```

## Binding an IPv6 Address to a MAC Hardware Address

Much like IPv4 addresses are bound to MAC hardware addresses in the ARP table, IPv6 addresses are bound to MAC hardware addresses in the neighbor discovery cache.

Use the **ipv6 neighbor** command in global configuration mode to statically bind an IPv6 address to a MAC hardware address.

The following example configures a static neighbor cache entry for IPv6 address **2001:11ac:fd34:3333:0:0:0:3** on a hardware device with a MAC address of **1111.1111.1111** on interface VLAN 51:

```
7100 Chassis(su)->configure
7100 Chassis(su-config)->ipv6 neighbor 2001:11ac:fd34:3333:0:0:0:3 1111.1111.1111 interface vlan.0.2501
7100 Chassis(su-config)->show ipv6 neighbor
FLAGS:    I = Incomplete      R = Reachable
          S = Stale           D = Delay
          P = Probe           L = Local
          F = Fixed (Static)  H = Host Owned
```

IPv6 Address	Hardware Address	Flg	Age	Updated	Expire	Interface	Port
2001:11ac:fd34:3333:0:0:0:3	11-11-11-11-11-11	FR		1m	-	- vlan.0.2501	host.0.1
2501:0:0:0:0:0:0:4	00-1f-45-5b-f5-cf	LR		1h01m	-	- vlan.0.2501	host.0.1
fe80:0:0:0:21f:45ff:fe5b:f5cf	00-1f-45-5b-f5-cf	LR		1h01m	-	- vlan.0.2501	host.0.1
2502:0:0:0:0:0:0:4	00-1f-45-5b-f5-cf	LR		18h26m	-	- vlan.0.2502	host.0.1
fe80:0:0:0:21f:45ff:fe5b:f5cf	00-1f-45-5b-f5-cf	LR		18h26m	-	- vlan.0.2502	host.0.1
fe80:0:0:0:21f:45ff:fe5b:f5cf	00-1f-45-5b-f5-cf	LR		18h26m	-	- vlan.0.2503	host.0.1
fe80:0:0:0:21f:45ff:fe5b:f5cf	00-1f-45-5b-f5-cf	LR		18h26m	-	- vlan.0.2504	host.0.1
3014:0:0:0:0:0:0:4	00-1f-45-5b-f5-cf	LR		18h26m	-	- vlan.0.3014	host.0.1
fe80:0:0:0:21f:45ff:fe5b:f5cf	00-1f-45-5b-f5-cf	LR		18h26m	-	- vlan.0.3014	host.0.1

```
Neighbor Entries Found: 9
7100 Chassis(su-config)->
```

## Configuring IPv6 Neighbor Discovery

[Procedure 29-3](#) describes how to configure a static IPv6 neighbor discovery cache entry.

**Procedure 29-3 Configuring an IPv6 Static Neighbor Discovery Cache Entry**

Task	Command(s)
In configuration command mode, optionally create a static binding between an IPv6 address to a MAC hardware address.	<b>ipv6 neighbor</b> <i>ipv6-address hardware-address interface interface</i>
In any command mode, verify the IPv6 neighbor discovery cache configuration.	<b>show ipv6 neighbors</b> [ <i>ipv6-address</i> ] [ <i>group</i> ] [ <i>interface interface</i> ] [ <i>verbose</i> ] [ <i>statistics</i> ]
In interface configuration mode, Optionally modify the number of neighbor discovery neighbor solicitation messages to send during duplicate address detection on unicast IPv6 addresses on the interface.	<b>ipv6 nd dad attempts</b> <i>num</i>

## The ARP Table

Address Resolution Protocol (ARP) is the method for finding a MAC hardware address when only the IP address is known. The 7100-Series firmware allows you to configure Address Resolution Protocol (ARP) table entries. ARP is used to associate IP addresses with MAC addresses. Once determined, the IP address and MAC association is stored in an ARP cache for rapid retrieval. An IP datagram is then encapsulated into a link-layer frame and sent over the network. A retransmit time period can be set to determine how often ARP requests are transmitted.

ARP table entries can be temporary or permanent. A temporary ARP entry has a timeout interval associated with it. The ARP entry expires at the end of the timeout interval. Expired ARP entries are referred to as stale entries. A stale entry timeout value determines how long the stale entry remains in the ARP table before it is removed.

Use the **set arp** command to configure a permanent ARP entry with the option of setting the entry to temporary.

Use the **show arp** command to display ARP table entries.

Use the **clear arp** command to clear specific static or all temporary ARP entries from the ARP table.

## Gratuitous ARP

Gratuitous ARP provides an ARP announcement packet containing valid sender hardware and protocol addresses for the host that sent it. Such a request is not intended to solicit a reply, but merely updates the ARP caches of other hosts that receive the packet. Gratuitous ARP is usually formatted as an ARP request, but it may also be formatted as an ARP reply. ARP announcements are typically sent out during startup. This helps to resolve problems which would otherwise occur if, for example, an IP-address-to-MAC-address mapping changed because a network card was replaced. In this example, gratuitous ARP solves the problem of remote hosts that still have the old mapping in their ARP caches. The 7100-Series provides for setting the behavior when an ARP announcement is received for an already existing ARP table entry or for a non-existing ARP table entry, referred to as ARP learning.

IP gratuitous ARP is enabled by default for the modification of pre-existing ARP table entries and is disabled by default for the learning of new table entries.

Use the **ip gratuitous-arp** command in interface configuration command mode to:

- Configure the device to ignore gratuitous ARP announcements received for existing ARP table entries

- Configure the device to change the ARP table only if the gratuitous ARP announcement is a reply
- Configure the device to change the ARP table only if the gratuitous ARP announcement is a request.

Use the **ip gratuitous-arp-learning** command, in interface configuration command mode, to allow an interface to learn new ARP bindings using gratuitous ARP. The router will learn new ARP bindings from reply, request, or both types of ARP announcements, based upon the option specified in this command.

Gratuitous ARP configuration does not affect normal ARP operations. Normal ARP packets (non gratuitous) will always be learned and updated regardless of gratuitous ARP configuration.

## Proxy ARP

Proxy ARP provides for the ability of a device on a given network to answer the ARP queries for a network address that is not on that network. The ARP Proxy, being aware of the traffic destination's location, provides its own MAC address in reply. Serving as an ARP Proxy for another host effectively directs LAN traffic to the Proxy. The "directed" traffic is then typically routed by the proxy to the intended destination via another interface.

Proxy ARP is enabled by default. Typically, proxy arp is only used to reply to requests for hosts that are reachable via a non-default route. Proxy ARP can be configured to respond to ARP requests for hosts that are only reachable via the default route. Proxy ARP can also be configured to respond to ARP requests that are received on the interface to which this command is applied, if the source IP address of the request is reachable on the local interface.

## ARP/ND Proxy-All

In an effort to reduce the amount of broadcast and multicast traffic in a broadcast domain, the router can be configured to respond to all ARP and Neighbor Discovery requests. The router will respond with the hardware address of the host that owns the IP address, if the ARP/ND entry for the Target Address exists in the router's ARP/ND cache. If the Target Address does not exist in the ARP/ND cache or if the entry is unresolved, the ARP/ND request will be dropped.

Use the command **arp-nd-proxy-all** in interface configuration mode to enable ARP/ND proxy all.

## Removing the Multicast ARP Restriction

As specified in RFC 1812, by default the router must not believe any ARP packet that claims the packet MAC address is broadcast or multicast. The multicast restriction can be removed on the interface using the **ip multicast-arp-learning** command in interface configuration mode.

## ARP Configuration Examples

The following example:

- Temporarily configures the IP address 10.21.128.1, MAC address 00:00:5e:00:01:01 binding in the ARP table

```
7100 Chassis(rw)->set arp 10.21.128.1 00:00:5e:00:01:01 temp
```

```
7100 Chassis(rw)->show arp
```

```
FLAGS:      U = Unresolved      S = Static
            L = Local        V = VRRP
            * = Stale       B = Best Guess Interface
```

IP Address	Hardware Address	Flg	Age	Updated	Interface	Port
10.21.128.1	00:00:5e:00:01:01	B	4h55m	1m	vlan.0.1	ge.1.1
10.21.130.59	00:11:88:0c:9f:78	L	5h05m	-	vlan.0.1	host.0.1

ARP Entries Found: 2

7100 Chassis(rw-config)->

The following example enables gratuitous ARP and ARP learning for ARP replies on VLAN 1:

```
7100 Chassis(rw)->configure
7100 Chassis(rw-config)->interface vlan 1
7100 Chassis(rw-config-intf-vlan.0.1)->ip gratuitous-arp reply
7100 Chassis(rw-config-intf-vlan.0.1)->ip gratuitous-arp-learning reply
7100 Chassis(rw-config-intf-vlan.0.1)->exit
7100 Chassis(rw-config)->
```

The following example enables proxy ARP for both default and local routes:

```
7100 Chassis(rw)->configure
7100 Chassis(rw-config)->ip prox
7100 Chassis(rw-config)->interface vlan 1
7100 Chassis(rw-config-intf-vlan.0.1)->ip proxy-arp default-route local
7100 Chassis(rw-config-intf-vlan.0.1)->exit
7100 Chassis(rw-config)->
```

[Procedure 29-4](#) describes how to configure the ARP table.

#### Procedure 29-4 Configuring the ARP Table

Step	Task	Command(s)
1.	Add mapping entries to the ARP table with the option of configuring them as temporary.	<b>set arp</b> <i>ip-address mac-address</i> [ <b>interface interface</b> ] [ <b>temp</b> ]
2.	Optionally, in interface configuration command mode, override the default ARP update process. <ul style="list-style-type: none"> <li>• <b>ignore</b> - Ignore all gratuitous ARP frames, no updates will occur. This option will also prevent any new learning from gratuitous arps, if the command <code>ip gratuitous-arp-learning</code> was used.</li> <li>• <b>reply</b> - Update from gratuitous arp reply only.</li> <li>• <b>request</b> - Update from gratuitous arp request only.</li> </ul>	<b>ip gratuitous-arp</b> { <b>ignore</b>   <b>reply</b>   <b>request</b> }

**Procedure 29-4 Configuring the ARP Table (continued)**

Step	Task	Command(s)
3.	<p>Optionally, in interface configuration command mode, allow an interface to learn new ARP bindings using gratuitous ARP.</p> <ul style="list-style-type: none"> <li>• <b>both</b> - Allows learning from both gratuitous arp reply and request.</li> <li>• <b>reply</b> - Allows learning from gratuitous arp reply.</li> <li>• <b>request</b> - Allows learning from gratuitous arp request.</li> </ul>	<b>ip gratuitous-arp-learning {both   reply   request}</b>
4.	<p>Optionally, in interface configuration command mode, enable proxy ARP on an interface.</p> <ul style="list-style-type: none"> <li>• <b>default-route</b> - Sets the router to respond to ARP requests for hosts that are only reachable via the default route. Typically, proxy arp is only used to reply to requests for host that are reachable via a non-default route.</li> <li>• <b>local</b> - Allows the router to respond to ARP requests that are received on the interface to which this command is applied if the target IP address of the request is reachable on the interface that received the request.</li> </ul>	<b>ip proxy-arp [default-route] [local]</b>
5.	<p>Optionally, in interface configuration command mode, remove the multicast restriction on ARP packets.</p>	<b>ip multicast-arp-learning</b>
6.	<p>Optionally, in interface configuration mode, configure the router to respond to all ARP and Neighbor Discovery requests.</p>	<b>arp-nd-proxy-all</b>

## Router Management and Information Display

[Table 29-2](#) lists routing parameters and their default values.

**Table 29-2 Default IP Routing Parameters**

Parameter	Description	Default Value
ARP entry type	Specifies whether an ARP table entry is permanent or temporary.	permanent
ARP retransmit time	Specifies the duration in seconds to wait before retransmitting ARP requests.	1 second
ARP stale entry timeout	Specifies the duration in seconds an ARP entry will remain in the stale state before the entry is removed from the ARP table.	1200 seconds

**Table 29-2 Default IP Routing Parameters (continued)**

Parameter	Description	Default Value
global router	Specifies the default router used when configuring the router directly from configuration command mode. The current implementation supports a single global router	global
gratuitous ARP	A feature that overrides the normal ARP updating process by providing an ARP announcement packet containing valid sender hardware and protocol addresses for the host that sent it.	enabled for ARP replies and ARP requests
gratuitous ARP learning	A feature that allows an interface to learn new ARP bindings using gratuitous ARP.	disabled
IPv4 forwarding	Specifies whether or not the routing interface will forward IPv4 traffic.	disabled
IPv6 address autoconfiguration	Specifies whether IPv6 addresses are auto configured on the interface.	disabled
IPv6 forwarding	Specifies whether or not the routing interface will forward IPv6 traffic.	disabled
neighbor discovery Duplicate Address Detection (DAD)	Specifies the number of DAD messages neighbor discovery will send out to attempt to determine whether the “tentative” address for this interface is a duplicate of another address in the network.	1 attempt
proxy ARP	A feature that provides for the ability of a device on a given network to answer the ARP queries for a network address that is not on that network.	enabled (no local or default-route)

[Table 29-3](#) describes how to manage IP configuration.

**Table 29-3 Managing the Router**

Task	Command
To clear this router configuration:	<b>clear router vrf</b> <i>vrf-name</i>
To delete one or all entries from the ARP table:	<b>clear arp</b> [ <i>ip-address</i>   <b>all</b> ]
To delete all non-static (dynamic) entries from the ARP table:	<b>clear arp-cache</b> [ <i>ip-address</i> ] [ <b>interface</b> <i>interface</i> ]

[Table 29-4](#) describes how to display IP configuration information and statistics.

**Table 29-4 Displaying IP Routing Information and Statistics**

Task	Command
To display router configuration:	<b>show router</b> [ <i>name</i> ]
To display the application limits for this router:	<b>show limits</b> [ <b>application</b> <i>application</i> ]



**Table 29-4 Displaying IP Routing Information and Statistics (continued)**

Task	Command
To display non-default, user entered configuration, or all configuration for this router:  Supported applications can be determined by entering the <b>show running-config ?</b> command.	<b>show running-config</b> [all] [application [all]]
To display configuration information for one or more interfaces:	<b>show interface</b> [interface-name]
To display configuration information for one or more IPv4 routing interfaces:	<b>show ip interface</b> [interface-name] [brief]
To display configuration information for one or more IPv6 routing interfaces:	<b>show ipv6 interface</b> [interface-name [prefix]] [brief]
To display information about IP protocols running on this device:	<b>show ip protocols</b>
To display information about IP routes:	<b>show ip route</b> [host [connected   host-address   dynamic   static]] [dest-address [prefix-mask]   prefix/prefix-length   connected   summary]
To display the device's ARP table:	<b>show arp</b> [ip-address] [interface interface] [statistics]
To display debug IP packet utility settings:	<b>show debugging</b>

## Terms and Definitions

[Table 29-5](#) lists terms and definitions used in this IP routing configuration discussion.

**Table 29-5 IP Routing Terms and Definitions**

Term	Definition
Address Resolution Protocol (ARP)	A protocol providing a method for finding a MAC hardware address when only the IP address is known.
ARP proxy	Provides for the ability of a device on a given network to answer the ARP queries for a network address that is not on that network.
blackhole route	Silently drops packets destined for this route's subnet.
Duplicate Address Detection (DAD)	An IPv6 neighbor discovery capability that uses neighbor solicitation and neighbor advertisement messages to verify the uniqueness of an address.
general prefix	The ability to assign a name to represent a network prefix from which longer IPv6 addresses can be configured.
global router	The default router from which VRF routing instances are configurable.
gratuitous ARP	A method for overriding the normal ARP process that provides an ARP announcement packet containing valid sender hardware and protocol addresses for the host that sent it. ARP announcements are sent out during startup.
IP address	An address used by the IP protocol to identify a routing interface or routing device.
managed address configuration	A DHCPv6 capability that determines whether the interface will send out IPv6 address configuration to an interface with IPv6 autoconfiguration enabled.

**Table 29-5 IP Routing Terms and Definitions (continued)**

<b>Term</b>	<b>Definition</b>
management interface	A non-forwarding interface to which an IP subnet can be assigned, allowing the network administrator to create an out-of-band management subnet designed to only pass network management data.
neighbor discovery	An IPv6 protocol defined in RFC4861 that uses ICMPv6 messages to determine the link-layer addresses of nodes residing on the same local link, to locate neighboring routers, to learn certain link and address configuration, and to track the reachability of neighbors.
neighbor unreachability detection	An IPv6 neighbor discovery capability that detects the failure of a neighbor or the failure of the forward path to the neighbor.
routing interface	A VLAN or loopback interface configured for IP routing.
static route	An administratively configured IP route consisting of the destination and next-hop IP addresses from the IP router the route is configured on.



## Routing Information Protocol (RIP) Configuration

This document describes the RIP feature and its configuration on Extreme Networks 7100-Series devices.

For information about...	Refer to page...
<a href="#">Using RIP in Your Network</a>	30-1
<a href="#">RIP Overview</a>	30-1
<a href="#">Configuring RIP</a>	30-4
<a href="#">Terms and Definitions</a>	30-5

### Using RIP in Your Network

The 7100-Series device supports Routing Information Protocol (RIP) Version 2. RIP is a distance-vector routing protocol for use in small networks; it is not intended for complex networks. RIP is described in RFC 2453. A router, running RIP broadcasts, updates at set intervals. Each update contains paired values where each pair consists of an IP network address and an integer distance to that network. RIP uses a hop count metric to measure the distance to a destination and is not appropriate for situations where routes need to be chosen based on real-time parameters such as a measured delay, reliability, or load.

The 7100-Series device implements plain text and MD5 authentication methods for RIP Version 2.

### RIP Overview

This section provides an overview of RIP configuration.

Enabling RIP on the device starts the RIP process which then begins populating its routing table and sending and receiving routing updates. Use the **router rip** command in configuration command mode to both enable RIP on the device and enter RIP configuration command mode.

Within RIP configuration command mode:

- Attach one or more networks to the RIP process specifying the IP address of the directly connected network, followed by the wildcard mask for this network. RIP network wildcard masks are reverse networks (use 1's for don't care bits). Use the **network** command to attach one or more networks to this RIP process.
- Optionally change the preference value for using RIP as the routing protocol for this device by changing the RIP administrative distance value using the **distance** command.
- Optionally specify interfaces which will not transmit any RIP update packets using the **passive-interface** command.

- Optionally adjust routing timers associated with:
  - The frequency of routing updates by specifying the interval, in seconds, at which routing updates are sent
  - The expiration of routes by specifying the interval, in seconds, from the point of the last update after which a route times out and is marked as expired
  - The deletion of routes by specifying the interval in seconds from the point of a routes expiration after which a route is deleted from the routing table

Using the **timers** command. Use the **show ip protocols** command to display RIP timer values.

- Specify route types that can be redistributed in RIP update messages using the **redistribute** command. The 7100-Series supports redistribution of connected and static routes, optionally specifying the hop count metric for these routes, or specifying OSPF using the process ID to redistribute.

### RIP Configuration Example

The following configuration example:

- Enables the RIP process on this device and enters RIP configuration command mode
- Attaches the 10.10.20.0 and 10.10.50.0 networks to this RIP process
- Configures VLANs 10 and 20 as passive-interfaces
- Changes the RIP timers to a 25 second update time, a 150 second expiration interval, and a 100 second flush time:
- Configures the redistribution of OSPF process ID 16546 routes over this RIP process

```
7100 Chassis(rw-config)->router rip
7100 Chassis(rw-config-rip)->network 10.10.20.0 0.0.0.255
7100 Chassis(rw-config-rip)->network 10.10.50.0 0.0.0.255
7100 Chassis(rw-config-rip)->passive-interface vlan 10
7100 Chassis(rw-config-rip)->passive interface vlan 20
7100 Chassis(rw-config-rip)->timers basic 25 150 100
7100 Chassis(rw-config-rip)->redistribute ospf 16546
7100 Chassis(rw-config-rip)->exit
7100 Chassis(rw-config-)->
```

## Configuring RIP Authentication

At the interface command level, RIP supports authentication configuration.

The authentication mode applied to the interface can be either clear text or encrypted MD5. Use the **ip rip authentication** mode command to specify the authentication mode for this interface.

Authentication parameters are specified in a key chain. The key chain can be configured for up to 255 keys. A key contains the key authentication string that is sent and received in RIP packets, an accept-lifetime that specifies the period during which an authentication key is valid to be received, and a send-lifetime which specifies the time period during which an authentication key is valid to be sent.

Use the **key chain** command in configuration command mode to enter key chain configuration command mode.

Use the **key** command in key chain configuration command mode to configure a key chain key and enter key configuration command mode.

Use the **key-string** command in key configuration command mode to specify the key string associated with this key.

Use the **accept-lifetime** command in key configuration command mode to specify the time period during which this key can be received for authentication by interface this key chain is associated with.

Use the **send-lifetime** command in key configuration command mode to specify the time period during which this key can be sent by the interface this key chain is associated with.

Use the **ip rip authentication keychain** command in interface configuration command mode to specify the named key chain this interface will use when authenticating RIP packets.

The following example:

- configures key **3** on key chain **md5key**, with a key string of **password**, an accept-lifetime and send-lifetime from the current time to infinite
- Configures VLAN 5 for RIP MD5 authentication
- Applies the md5key key chain to VLAN 5

```

7100 Chassis(rw-config)->key chain md5key
7100 Chassis(rw-config-keychain)->key 3
7100 Chassis(rw-config-keychain-key)->key-string password
7100 Chassis>Router(config-keychain-key)->accept-lifetime 02:30:00 jul 30 2009
infinite
7100 Chassis(rw-config-keychain-key)->send-lifetime 02:30:00 jul 30 2009 infinite
7100 Chassis(rw-config-keychain-key)->show running config
.
.
.
!
key chain md5key
  key 3
  key-string password
  accept-lifetime 02:30:00 Jul 30 2009 06:28:14 Feb 7 2106
  send-lifetime 02:30:00 Jul 30 2009 06:28:14 Feb 7 2106
  exit
exit
!
7100 Chassis(rw-config-keychain-key)->exit
7100 Chassis(rw-config-keychain)->exit
7100 Chassis(rw-config)->interface vlan 5
7100 Chassis(rw-config-intf-vlan.0.5)->ip rip authentication mode md5
7100 Chassis(rw-config-intf-vlan.0.5)->ip rip authentication keychain md5key
7100 Chassis(rw-config-intf-vlan.0.5)->exit
7100 Chassis(rw-config)->

```

## Configuring RIP Offset

In interface command mode, an offset can be added to the hop metric of an incoming or outgoing route learned by RIP. Use the **ip rip offset** command, specifying an offset value and whether the offset applies to incoming or outgoing route.

The following example configures VLAN 1 with a RIP offset of 5 for incoming RIP learned routes:

```
7100 Chassis(rw-config)->interface vlan 1
7100 Chassis(rw-config-intf-vlan.0.1)->ip rip offset in 1
```

## Configuring RIP

This section provides details for the configuration of RIP on the 7100-Series products.

[Table 30-1](#) lists RIP parameters and their default values.

**Table 30-1 Default RIP Parameters**

Parameter	Description	Default Value
RIP process	The RIP Router process on this device.	disabled
distance	The administrative distance that specifies the preference for RIP routing over other routing types on this device.	120
update interval	Specifies the interval between routing updates.	30 seconds
expiration interval	Specifies the interval from the point of the last update after which a route times out and is marked to expire.	180 seconds
flush interval	Specifies the interval from the point of a routes expiration after which a route is deleted from the routing table.	120 seconds

[Procedure 30-1](#) describes how to configure RIP.

**Procedure 30-1 Configuring RIP**

Step	Task	Command(s)
1.	In configuration command mode, enable the RIP process for this device.	<b>router rip</b>
2.	In RIP configuration command mode, attach one or more networks to this RIP process.	<b>network</b> <i>ip-address wild-card-bits</i>
3.	Optionally, in RIP configuration command mode, change the administrative distance for RIP routing on this device.	<b>distance</b> <i>weight</i>
4.	Optionally, in interface configuration command mode, add an offset to the hop metric of an incoming or outgoing RIP route for this interface.	<b>ip rip offset</b> { <i>in</i>   <i>out</i> } <i>value</i>

**Procedure 30-1 Configuring RIP (continued)**

Step	Task	Command(s)
5.	Optionally, in RIP configuration command mode, change the basic timers associated with RIP: <ul style="list-style-type: none"> <li>• Update interval</li> <li>• Route expiration interval</li> <li>• Route flush interval</li> </ul>	<b>timers basic</b> <i>update-seconds invalid-seconds flush-seconds</i>
6.	Optionally, in configuration command mode, name a RIP authentication key chain and enter key chain configuration command mode.	<b>key chain</b> <i>name</i>
7.	Optionally, in key chain configuration command mode, create a RIP authentication key for this key chain and enter authentication key configuration command mode.	<b>key</b> <i>key-id</i>
8.	Optionally, In authentication key configuration command mode, specify a key-string for this key that will be used by RIP to authenticate sent and received RIP packets.	<b>key-string</b> <i>text</i>
9.	Optionally, in key configuration command mode, specify a time period during which an authentication key is valid to be received.	<b>accept-lifetime</b> <i>start-time month date year {duration seconds   end-time   infinite}</i>
10.	Optionally, in key configuration command mode, specify a time period during which an authentication key is valid to be sent.	<b>send-lifetime</b> <i>start-time month date year {duration seconds   end-time   infinite}</i>
11.	Optionally, in interface configuration command mode, apply a RIP authentication key chain to an interface.	<b>ip rip authentication keychain</b> <i>name</i>
12.	Optionally, in interface configuration command mode, set the authentication mode when a key chain is present on this interface.	<b>ip rip authentication mode</b> <i>{text   md5}</i>
13.	Optionally, in RIP configuration command mode, specify an interface that will be prevented from transmitting RIP update packets.	<b>passive-interface vlan</b> <i>vlan-id</i>
14.	In RIP configuration command mode, specify the non-RIP protocols to be distributed in RIP update messages.	<b>redistribute</b> <i>{connected   ospf process-id   static} [metric metric-value]</i>

## Terms and Definitions

[Table 30-2](#) lists terms and definitions used in this RIP configuration discussion.

**Table 30-2 RIP Configuration Terms and Definitions**

Term	Definition
Routing Information Protocol (RIP)	A distance-vector routing protocol for use in small networks that broadcasts route updates at set intervals using a hop metric to determine route preference.
distance	An administrative value that sets the preference for the routing protocols on this device.



**Table 30-2 RIP Configuration Terms and Definitions (continued)**

<b>Term</b>	<b>Definition</b>
RIP offset	A value that is added to the hop metric of an incoming or outgoing route learned by RIP for the configured interface.
update interval	Sets the interval that determines the frequency of routing updates.
expiration interval	Sets the interval that determines the expiration of a route based upon the point of the last update.
flush interval	Sets the interval for the deletion of an expired route based upon the point of expiration.
key chain	A named chain that holds RIP authentication keys.
key	A key chain member that contains the key string used to authenticate RIP packets, accept-lifetime, and send-lifetime.
key string	A text string that is sent with RIP packets which must agree at both ends of the transmission for authentication to take place.
accept-lifetime	Specifies the time period during which an authentication key on a key chain is valid to be received by this device.
send-lifetime	Specifies the time period during which an authentication key on a key chain is valid to be sent by this device.
passive-interface	An interface configured to not transmit RIP update packets.

## Routing Information Protocol Next Generation (RIPng) Configuration

This document describes the RIPng feature and its configuration on Extreme Networks 7100-Series devices.

For information about...	Refer to page...
<a href="#">Using RIPng in Your Network</a>	31-1
<a href="#">RIPng Configuration Overview</a>	31-2
<a href="#">Configuring RIPng</a>	31-3
<a href="#">Terms and Definitions</a>	31-4

### Using RIPng in Your Network

The 7100-Series device supports Routing Information Protocol Next Generation (RIPng). RIPng is a distance-vector routing protocol for use in small networks; it is not intended for complex networks. RIPng is an Interior Gateway Protocol (IGP) in that it is used within a single Autonomous System (AS). RIPng is described in RFC 2080. A router, running RIPng broadcasts, updates at set intervals. Each update contains paired values where each pair consists of an IP network address and an integer distance to that network. RIPng uses a hop count metric to measure the distance to a destination and is not appropriate for situations where routes need to be chosen based on real-time parameters such as a measured delay, reliability, or load.

RIPng is conceptually the same as RIPv2 for IPv4. In essence, the IPv4 address is expanded into an IPv6 address. RIPng replaces the IPv4 subnet with an IPv6 prefix length. The next-hop header field is eliminated but the functionality is preserved. The route tag field is preserved. The maximum diameter (metric value) of the network is 15, assuming that a cost of 1 is used for each network; 16 still means the route is unreachable.

RIPng uses fixed metrics to compare alternative routes. The RIPng metric of a network is an integer value range of 1 - 15. Given that the maximum path limit is 15, the metric value is usually set to 1.

RIPng uses a UDP-based protocol and sends and receives datagrams on UDP port 521, which is used for all communications with another router's RIPng process. RIPng messages are either a request for all or a part of the receiving router's route table or a response that contains all or part of a sending router's route table.

Unsolicited response messages containing the complete routing table of the sending router are sent out every 30 seconds by default to every neighboring router. There are two timers associated with routing table entries. The expiration timer specifies when a route has expired. The expiration timer is initialized when a route is first established and each time an update is received for that route. The flush timer specifies when an expired route should be removed from the route table.

Each network has an IPv6 destination address prefix and prefix length associated with it.

Authentication has been removed from RIPng.

## RIPng Configuration Overview

This section provides an overview of RIPng configuration.

Enabling RIPng on the device starts the RIPng process which then begins populating its routing table and sending and receiving routing updates. Use the **ipv6 router rip** command in configuration command mode to both enable RIPng on the device and enter RIPng configuration command mode.

Within RIPng configuration command mode:

- Optionally change the preference value for using RIPng as the routing protocol for this device by changing the RIPng administrative distance value using the **distance** command.
- Optionally assign standard ACLs to a distribution list to filter networks received and to suppress networks from being advertised in RIPng updates. Use the **distribute-list** command to assign an ACL to the list, specifying an IPv6 routing configured VLAN and whether the affected routing updates are incoming or outgoing.
- Optionally specify interfaces which will not transmit any RIPng update packets using the **passive-interface** command.
- Specify route types that can be redistributed in RIPng update messages using the **redistribute** command. The 7100-Series supports redistribution of connected, OSPF (specifying the process ID), and static routes, optionally specifying the hop count metric for these routes or a route map.
- Optionally, using the **timers** command, adjust routing timers associated with:
  - The frequency of routing updates by specifying the interval, in seconds, at which routing updates are sent
  - The expiration of routes by specifying the interval, in seconds, from the point of the last update after which a route times out and is marked as expired
  - The deletion of routes by specifying the interval in seconds from the point of a routes expiration after which a route is deleted from the routing table

Because IPv6 addressing prefix is ambiguous concerning network addressing compared to IPv4, you do not directly specify a network destination address. Instead, you enable RIPng on the interface using the **ipv6 rip enable** command and the interface address is automatically set as the network destination address.

An offset can be added or removed to the hop metric for all incoming or outgoing RIPng routes for a given interface. Adding an offset is used for the purpose of making an interface a backup. Use the **ipv6 rip offset** command in interface configuration mode to add or remove an offset for either incoming or outgoing RIPng routes.

### RIPng Configuration Example

The following configuration example:

- Enables the RIPng process on this device and enters RIPng configuration command mode
- Configures VLANs **10** as a passive-interface
- Changes the RIPng timers to a **25** second update time, a **150** second expiration interval, and a **100** second flush time:
- Configures the redistribution of OSPF process ID **16546** routes over this RIPng process

- Configures IPv6 access-list **ipv6list1** with a rule to deny route **2001:0:0:0:21f:45ff:fe3d:21be/64** and applies the ACL to the distribution-list for outgoing packets on VLAN 20
- Enables RIPng on VLANs 10 and 20

```

7100 Chassis(rw-config)->ipv6 router rip
7100 Chassis(rw-config-ripng)->passive-interface vlan 10
7100 Chassis(rw-config-ripng)->timers basic 25 150 100
7100 Chassis(rw-config-ripng)->redistribute ospf 16546
7100 Chassis(rw-config-ripng)->exit
7100 Chassis(rw-config)->ipv6 access-list standard ipv6list1
7100 Chassis(rw-cfg-ipv6-std-acl)->deny 2001:0:0:0:21f:45ff:fe3d:21be/64
7100 Chassis(rw-cfg-ipv6-std-acl)->exit
7100 Chassis(rw-config)->ipv6 router rip
7100 Chassis(rw-config-ripng)->distribute-list ipv6list1 out vlan 20
7100 Chassis(rw-config-ripng)->exit
7100 Chassis(rw-config)->interface vlan 10
7100 Chassis(rw-config-intf-vlan.0.10)->ipv6 rip enable
7100 Chassis(rw-config-intf-vlan.0.10)->exit
7100 Chassis(rw-config)->interface vlan 20
7100 Chassis(rw-config-intf-vlan.0.20)->ipv6 rip enable
7100 Chassis(su-config-intf-vlan.0.20)->exit
7100 Chassis(rw-config)->

```

## Configuring RIPng

This section provides details for the configuration of RIPng on the 7100-Series products.

[Table 31-1](#) lists RIPng parameters and their default values.

**Table 31-1 Default RIPng Parameters**

Parameter	Description	Default Value
RIPng process	The RIPng Router process on this device.	disabled
distance	The administrative distance that specifies the preference for RIPng routing over other routing types on this device.	120
update interval	Specifies the interval between routing updates.	30 seconds
expiration interval	Specifies the interval from the point of the last update after which a route times out and is marked to expire.	180 seconds
flush interval	Specifies the interval from the point of a routes expiration after which a route is deleted from the routing table.	120 seconds

[Procedure 31-1](#) describes how to configure RIPng.

### Procedure 31-1 Configuring RIPng

Step	Task	Command(s)
1.	In configuration command mode, enable the RIPng process for this device.	<b>ipv6 router rip</b>
2.	Optionally, in RIPng configuration command mode, change the administrative distance for RIPng routing on this device.	<b>distance <i>weight</i></b>
3.	Optionally, in RIPng configuration command mode, change the basic timers associated with RIPng: <ul style="list-style-type: none"> <li>• Update interval</li> <li>• Route expiration interval</li> <li>• Route flush interval</li> </ul>	<b>timers basic <i>update-seconds expiration-seconds flush-seconds</i></b>
4.	Optionally, in RIPng configuration command mode, specify an interface that will be prevented from transmitting RIPng update packets.	<b>passive-interface vlan <i>vlan-id</i></b>
5.	Optionally, in RIPng configuration command mode, specify a standard IPv6 ACL to be added to the distribute-list to filter networks received and to suppress networks from being advertised in RIPng updates.	<b>distribute-list <i>access-list-name</i> {in vlan <i>vlan-id</i>   out vlan <i>vlan-id</i>}</b>
6.	In RIPng configuration command mode, specify the non-RIPng protocols to be distributed in RIPng update messages.	<b>redistribute {connected   ospf <i>process-id</i>   static} [metric <i>metric-value</i>] [route-map <i>route-map</i>]</b>
7.	In interface configuration mode, enable RIPng on all interfaces that will use the protocol.	<b>ipv6 rip enable</b>
8.	Optionally, in interface configuration command mode, add an offset to the hop metric of an incoming or outgoing RIPng route for this interface.	<b>ipv6 rip offset {in   out} <i>value</i></b>

## Terms and Definitions

[Table 31-2](#) lists terms and definitions used in this RIPng configuration discussion.

**Table 31-2 RIPng Configuration Terms and Definitions**

Term	Definition
Routing Information Protocol Next Generation (RIPng)	A distance-vector routing protocol for use in small IPv6 networks that broadcast route updates at set intervals using a hop metric to determine route preference.
distance	An administrative value that sets the preference for the routing protocols on this device.
RIPng offset	A value that is added to the hop metric of an incoming or outgoing route learned by RIPng for the configured interface.
update interval	Sets the interval that determines the frequency of routing updates.

**Table 31-2 RIPng Configuration Terms and Definitions (continued)**

Term	Definition
expiration interval	Sets the interval that determines the expiration of a route based upon the point of the last update.
flush interval	Sets the interval for the deletion of an expired route based upon the point of expiration.
passive-interface	An interface configured to not transmit RIPng update packets.



## Open Shortest Path First (OSPFv2) Configuration

This chapter provides the following information about configuring and monitoring OSPFv2 on Extreme Networks 7100-Series devices:

For information about...	Refer to page...
<a href="#">Using the OSPF Protocol in Your Network</a>	<a href="#">32-1</a>
<a href="#">Implementing OSPF</a>	<a href="#">32-2</a>
<a href="#">OSPF Overview</a>	<a href="#">32-3</a>
<a href="#">Configuring OSPF</a>	<a href="#">32-22</a>



**Note:** OSPFv2 requires an advanced router license. See “[License Overview](#)” on page 9-10 for 7100-Series advanced router license details.

### Using the OSPF Protocol in Your Network

The Open Shortest Path First (OSPF) Link-state routing protocol is considered a TCP/IP internet routing Interior Gateway Protocol (IGP). OSPF distributes routing information between routers belonging to a single Autonomous System (AS). The OSPF protocol is based on link-state or SPF technology. The advantages associated with a link-state routing protocol are:

- Rapid convergence
- Reduced routing updates traffic over traditional distance-vector protocols

This OSPF implementation supports RFC 2328 *OSPF Version 2*.

The OSPF protocol is designed expressly for the TCP/IP internet environment. It provides for the authentication of routing updates, and utilizes IP multicast when sending and receiving the updates.

OSPF routes IP packets based solely on the destination IP address found in the IP packet header. IP packets are not encapsulated in any further protocol headers as they transit the Autonomous System. OSPF is a dynamic routing protocol in that it quickly detects topological changes in the AS, such as router interface failures, and calculates new loop-free routes after a period of convergence. This period of convergence is short and involves a minimum of routing traffic. In a link-state routing protocol, each router maintains a database describing the AS's topology. This database is referred to as the link-state database. Each participating router has an identical database. Each individual piece of this database is a particular router's local state made up of such information as the router's usable interfaces and reachable neighbors. The router distributes its local state throughout the AS by flooding.

Each network that has at least two attached routers has a designated router. The designated router generates an LSA for the network and has other special responsibilities in the running of the



protocol, enabling a reduction in the number of adjacencies required on a network. This in turn reduces the amount of routing protocol traffic and the size of the link-state database.

All routers run the exact same algorithm, in parallel. From the link-state database, each router constructs a tree of shortest paths with itself as root. This shortest-path tree provides the route to each destination in the AS. Externally derived routing information appears on the tree as leaves. When several equal-cost routes to a destination exist, traffic is distributed equally among them. The cost of a route is described by a single dimensionless metric.

OSPF allows sets of networks to be grouped together. Such a grouping is called an area. The topology of an area is hidden from the rest of the AS. This information hiding enables a significant reduction in routing traffic. Also, routing within the area is determined only by the area's own topology, lending the area protection against bad routing data. An area is a generalization of an IP subnetted network. OSPF enables the flexible configuration of IP subnets. Each route distributed by OSPF has a destination and mask. Two different subnets of the same IP network number may have different masks providing a different range of addresses for that subnet. This is commonly referred to as Variable Length Subnet Masking (VLSM). A packet is routed to the longest or most specific match. Host routes are considered to be subnets whose masks are "all ones" (0xffffffff).

All OSPF protocol exchanges are authenticated. This means that only trusted routers can participate in the AS's routing. The 7100-Series platform supports either simple or MD5 authentication schemes. Separate authentication schemes can be configured for each IP subnet.

Route redistribution is supported for RIP, connected, BGP, and static routes.

The Bidirectional Forwarding Detection (BFD) protocol providing sub-second failure detection on OSPF forwarding interfaces is enabled by default on all OSPF interfaces.

An OSPF Customer Edge (CE) router can be configured as a peer to a Provider Edge (PE) router by enabling the PE-CE protocol on the PE-CE associated routers.

## Implementing OSPF

To implement OSPF in your network:

- Map out the AS including routers, network subnets, and the areas to which they belong
- Configure each routing interface on each router with an IP address and mask
- Create an OSPF routing instance for this AS
- Configure the network addresses, masks, and areas for each router in the AS
- Configure each router with a router ID
- Optionally determine which router will be the designated router and backup and configure OSPF priority values accordingly
- Optionally configure OSPF timers
- Optionally, configure the protocols and route types that will be redistributed over this AS
- Optionally configure interface cost
- Optionally modify the administrative distance for OSPF routes
- Optionally configure either simple or MD5 authentication per interface
- Optionally configure areas including virtual-links, stub, and NSSA
- Optionally enable graceful-restart
- Optionally enable the BFD protocol on all OSPF interfaces
- Optionally enable the PE-CE protocol on the router, with PE-CE enabled

- Optionally, configure a domain tag for this router
- Optionally, configure a primary or secondary domain ID for this router
- Optionally redistribute BGP discovered routes over OSPF

## OSPF Overview

OSPF is enabled by creating an OSPF instance. Once an instance is created, the router's OSPF settings are configured with respect to the Instance ID and IP interfaces. By default, OSPF is disabled on the 7100-Series device. Be aware that unspecified parameters use their default values, and any parameters specified at the interface level will override the values specified at the area level.

### Configuring Basic OSPF Parameters

Basic OSPF configuration consists of:

- Entering interface configuration mode for the routing interfaces for this device
- Configuring each routing interface with an IP address and mask
- Enabling the interface
- Creating an OSPF routing instance
- Configuring the network address, mask, and area for this routing instance

### Configuring an IP Address

An IP address must be associated with any interface that will route traffic on the router. In interface configuration mode, configure the IP address for each routing interface using the **ip address** command specifying the IP address and mask. For example, IP address 10.10.10.1 would be specified as 10.10.10.1 255.255.255.0. Enable the interface using the **no shutdown** command.

### Configuring a Routing Instance

OSPF routing configuration takes place within a routing instance. Configure a routing instance using the **router ospf** command in global configuration command mode. Executing this command places you in the OSPF router configuration command mode for the specified OSPF router instance.

### Configuring Networks

A network is made up of a number of IP routers that belong to the same IP network, subnet, or supernet as determined by a device's combined IP address and mask. An edge connecting a router to a network indicates that the router has an interface on the network. Networks can be either transit or stub networks. Transit networks are those capable of carrying data traffic that is neither locally originated nor locally destined. A stub network has only incoming edges.

Use the **network** command in the OSPF router configuration command mode to configure networks and associated areas for this router. See section "[Configuring OSPF Areas](#)" on page 32-9 for information on OSPF areas and their configuration.



**Note:** OSPF network wildcard masks are reverse networks. This means that wherever there is a 1 in a regular netmask, use a 0 in a wildcard mask. For example, if the network mask is 255.255.255.0 (/24), specify a wildcard mask of **0.0.0.255**.

## Basic OSPF Topology

Figure 32-1 provides an overview of a basic OSPF topology. This topology displays two areas: a backbone area which must exist in any OSPF topology and a directly connected area 1. See “[Configuring OSPF Areas](#)” on page 32-9 for a full discussion of OSPF area configuration. This basic configuration requires the configuration of three interfaces and associated IP addresses, three networks, and two routers on a single OSPF router instance.

### Example


The following example configures the basic OSPF topology as displayed in [Figure 32-1](#) on page 32-5:

#### Router 1 CLI Input

```
Router 1 (rw) ->configure
Router 1 (rw-config) ->interface vlan 1
Router 1 (rw-config-intf-vlan.0.1) ->ip address 172.10.1.1 255.255.255.0
Router 1 (rw-config-intf-vlan.0.1) ->exit
Router 1 (rw-config) ->interface vlan 2
Router 1 (rw-config-intf-vlan.0.2) ->ip address 172.10.2.1 255.255.255.0
Router 1 (rw-config-intf-vlan.0.2) ->exit
Router 1 (rw-config) ->router ospf 1
Router 1 (rw-config-ospf-1) ->network 172.10.1.0 0.0.0.255 area 1
Router 1 (rw-config-ospf-1) ->network 172.10.2.0 0.0.0.255 area 1
Router 1 (rw-config-ospf-1) ->exit
Router 1 (rw-config) ->
```

#### Router 2 CLI Input

```
Router 2 (rw) ->configure
Router 2 (rw-config) ->interface vlan 2
Router 2 (rw-config-intf-vlan.0.1) ->ip address 172.10.2.2 255.255.255.0
Router 2 (rw-config-intf-vlan.0.1) ->exit
Router 2 (rw-config) ->interface vlan 3
Router 2 (rw-config-intf-vlan.0.2) ->ip address 172.2.1.1 255.255.255.0
Router 2 (rw-config-intf-vlan.0.2) ->exit
Router 2 (rw-config) ->router ospf 1
Router 2 (rw-config-ospf-1) ->network 172.10.2.0 0.0.0.255 area 1
Router 2 (rw-config-ospf-1) ->network 172.2.1.0 0.0.0.255 area 0
Router 2 (rw-config-ospf-1) ->exit
Router 2 (rw-config) ->
```

**Figure 32-1 Basic OSPF Topology**

## Configuring the Router ID

OSPF initially assigns all routers a router ID based on the highest loopback IP address of the interfaces configured for IP routing. If there is no loopback interface configured then it will be the highest VLAN IP address configured. This unique value, which is included in the hello packet transmitted in Link State Advertisements (LSA), identifies one router to another and helps establish adjacencies among OSPF routers. When you specify an interface as the router ID, this value supersedes the default ID.

### Example

The following example configures the router ID topology as displayed in [Figure 32-2](#) on page 32-6:

#### Router 1

```
Router 1 (rw) -> configure
Router 1 (rw-config) -> interface loopback 1
Router 1 (su-config-intf-loop.0.1) -> ip address 1.1.1.1 255.255.255.255
Router 1 (rw-config-intf-loop.0.1) -> exit
Router 1 (rw-config) -> router ospf 1
Router 1 (rw-config-ospf-1) -> network 10.1.2.2 0.0.0.255 area 1
Router 1 (rw-config-ospf-1) -> router-id 1.1.1.1
Router 1 (rw-config-ospf-1) -> exit
Router 1 (rw-config) ->
```

#### Router 2

```
Router 2 (rw) -> configure
Router 2 (rw-config) -> interface loopback 1
Router 2 (su-config-intf-loop.0.1) -> ip address 2.2.2.2 255.255.255.255
Router 2 (rw-config-intf-loop.0.1) -> exit
Router 2 (rw-config) -> router ospf 1
Router 2 (rw-config-ospf-1) -> network 10.1.2.1 0.0.0.255 area 1
```

```

Router 2 (rw-config-ospf-1)->network 10.2.3.1 0.0.0.255 area 0
Router 2 (rw-config-ospf-1)->router-id 2.2.2.2
Router 2 (rw-config-ospf-1)->exit
Router 2 (rw-config)->

```

### Router 3

```

Router 3 (rw)->configure
Router 3 (rw-config)->interface loopback 1
Router 3 (su-config-intf-loop.0.1)->ip address 3.3.3.3 255.255.255.255
Router 3 (rw-config-intf-vlan.0.1)->exit
Router 3 (rw-config)->router ospf 1
Router 3 (rw-config-ospf-1)->network 10.3.4.1 0.0.0.255 area 2
Router 3 (rw-config-ospf-1)->network 10.2.3.2 0.0.0.255 area 0
Router 3 (rw-config-ospf-1)->router-id 3.3.3.3
Router 3 (rw-config-ospf-1)->exit
Router 3 (rw-config)->

```


### Router 4

```

Router 4 (rw)->configure
Router 4 (rw-config)->interface loopback 1
Router 4 (su-config-intf-loop.0.1)->ip address 4.4.4.4 255.255.255.255
Router 4 (rw-config-intf-vlan.0.1)->exit
Router 4 (rw-config)->router ospf 1
Router 4 (rw-config-ospf-1)->network 10.3.4.2 0.0.0.255 area 2
Router 4 (rw-config-ospf-1)->router-id 4.4.4.4
Router 4 (rw-config-ospf-1)->exit
Router 4 (rw-config)->

```

**Figure 32-2 OSPF Router ID Topology**



## Configuring the Designated Router

In the process of implementing OSPF, a large number of multi-access links to routers across the network may cause too many adjacencies to form. To avoid this problem, a Designated Router (DR) is elected per multi-access network to build adjacencies to all other routers on that network.

A Backup Designated Router (BDR) is also elected in case the Designated Router (DR) fails, in which case the BDR will become the DR.



**Note:** A DR is required only for multi-access networks. Point-to-Point links do not need a DR because only a single adjacency is required.

To elect a DR from a host of candidates on the network, each router multicasts a hello packet and examines the priority of hello packets received from other routers. The router with the highest priority is elected the DR, and the router with the next highest priority is elected the BDR. Any router with a priority of 0 will opt out of the DR election process. See the “[Configuring Router Priority](#)” on page 32-7 for details on configuring router priority. If DR candidates all share non-zero priorities, OSPF applies the router ID as a tie-breaker where the highest ID is chosen DR and the next highest ID is chosen BDR.

## Configuring Router Priority

When two routers attached to a network both attempt to become the designated router, the one with the highest router priority takes precedence. A router whose router priority is set to 0 is ineligible to become the designated router on the attached network. Router priority is specified per router interface and is advertised in hello packets sent out by the interface.

Use the **ip ospf priority** command in interface configuration command mode to specify the router priority that will be specified for LSAs going out this interface. See “[Configuring the Designated Router](#)” on page 32-6 for a router priority configuration example.

[Figure 32-3](#) on page 32-8 displays a designated router topology example. The example will configure the four displayed routers with the following priorities:

- Router 1 = 25
- Router 2 = 10
- Router 3 = 30
- Router 4 = 0

Router 4 will not take part in the election process at all. Router 3 has the highest priority and therefore will be elected DR. Router 1 has the second highest priority and will be elected BDR.

## Example

The following example provides the input required to configure the designated router topology as displayed in [Figure 32-3](#) on page 32-8:

### Router 1

```
Router 1 (rw) ->configure
Router 1 (rw-config) ->interface vlan 1
Router 1 (rw-config-intf-vlan.0.1) ->ip ospf priority 25
Router 1 (rw-config-intf-vlan.0.1) ->exit
Router 1 (rw-config) ->
```

### Router 2

```
Router 2 (rw) ->configure
Router 2 (rw-config) ->interface vlan 1
Router 2 (rw-config-intf-vlan.0.1) ->ip ospf priority 10
Router 2 (rw-config-intf-vlan.0.1) ->exit
Router 2 (rw-config) ->
```


### Router 3

```
Router 3 (rw)->configure
Router 3 (rw-config)->interface vlan 1
Router 3 (rw-config-intf-vlan.0.1)->ip ospf priority 30
Router 3 (rw-config-intf-vlan.0.1)->exit
Router 3 (rw-config)->
```

### Router 4

```
Router 4 (rw)->configure
Router 4 (rw-config)->interface vlan 1
Router 4 (rw-config-intf-vlan.0.1)->ip ospf priority 0
Router 4 (rw-config-intf-vlan.0.1)->exit
Router 4 (rw-config)->
```

**Figure 32-3 OSPF Designated Router Topology**



## Configuring the Administrative Distance for OSPF Routes

If several routes coming from different protocols are presented to the Route Table Manager (RTM), the protocol with the lowest administrative distance will be chosen for route installation. The 7100-Series platform supports connected, static, OSPF, and RIP routes.

The table below displays the default distance for these routing protocols.

Route Source	Default Distance
Connected	0
Static	1
OSPF	110
RIP	120

Use the **distance ospf** command in OSPF router configuration command mode to change the administrative distance assigned to the OSPF protocol. This command provides for the configuration of separate values for OSPF external and intra-area routes.

## Configuring OSPF Areas

OSPF allows collections of contiguous networks and hosts to be grouped together. Such a group, together with the routers having interfaces to any one of the included networks, is called an area. Each area runs a separate copy of the basic link-state routing algorithm. This means that each area has its own link-state database.

The topology of an area is invisible from the outside of the area, and routers internal to a given area know nothing of the detailed topology external to the area. This isolation of area detail enables the protocol to effect a marked reduction in routing traffic as compared to treating the entire Autonomous System as a single link-state domain. A router has a separate link-state database for each area it is connected to. Routers connected to multiple areas are called Area Border Routers (ABR). Two routers belonging to the same area have, for that area, identical area link-state databases.

An autonomous system can have one or more areas. A multiple area AS must define one of the areas as the backbone with an area ID of 0. Area IDs are assigned during network configuration using the **network** command (see “[Configuring Networks](#)” on page 32-3). All non-backbone areas in a multiple area AS must either be contiguous to the backbone or connected using a virtual-link. The backbone is responsible for distributing routing information between non-backbone areas. The backbone must be contiguous. However, it need not be physically contiguous; backbone connectivity can be established and maintained through the configuration of virtual links.

Virtual links can be configured between any two backbone routers that have an interface to a common non-backbone area. Such virtual links belong to the backbone. The protocol treats two routers joined by a virtual link as if they were connected by an unnumbered point-to-point backbone network.

See RFC 2328 *OSPF Version 2* for further details on inter-area connectivity.

An Area ID can be any value from 0 - 4294967295, but is converted into the 32-bit dotted-quad format (area 50 would be displayed as 0.0.0.50; area 3546 would be displayed as 0.0.13.218)

## Configuring Area Range

An area range is a form of address summarization that defines a range of addresses to be used by the backbone ABRs when they communicate routes to other areas. Area range is a critical tool that pares the route tables and update traffic, as well as reduces network recalculation by the Dijkstra algorithm. Area range configuration summarizes by aggregating an areas’ internal networks to advertise a single network. Backbone routers see only one update, representing an entire range of subnets. Area ranges can be configured for purposes of network advertisement as well as summarization of subnets that should not be advertised.

Use the **area range** command in OSPF configuration command mode to configure an area network summarization.

## Example

The following example provides the input required to configure summarization of the three area topology as displayed in [Figure 32-4](#) on page 32-10:

### Area 1

```
ABR1 (rw) ->configure
ABR1 (rw-config) ->router ospf 1
```



```

ABR1 (rw-config-ospf-1) -> area 1 range 10.2.0.0 255.255.0.0
ABR1 (rw-config-ospf-1) -> exit
ABR1 (rw-config) ->
    
```

### Area 2

```


ABR2 (rw) -> configure
ABR2 (rw-config) -> router ospf 1
ABR2 (rw-config-ospf-1) -> area 2 range 10.3.0.0 255.255.0.0
ABR2 (rw-config-ospf-1) -> area 2 range 10.3.2.0 255.255.255.0 not-advertised
ABR2 (rw-config-ospf-1) -> exit
ABR2 (rw-config) ->
    
```

### Area 3

```

ABR3 (rw) -> configure
ABR3 (rw-config) -> router ospf 1
ABR3 (rw-config-ospf-1) -> area 3 range 10.1.0.0 255.255.0.0
ABR3 (rw-config-ospf-1) -> exit
ABR3 (rw-config) ->
    
```

**Figure 32-4 OSPF Summarization Topology**



### Configuring a Stub Area

A stub area is a non-transit area. In other words, an area that does not originate or propagate external routes. AS-external-LSAs are not flooded into the stub area; routing to AS external

networks is based on a single per-area default route. This reduces the link-state-database size and memory requirements for routers within stub areas.

Handy for reducing routing table size, a stub area is a “dead-end” in which there is no other way to enter or exit except through an Area Border Router (ABR). No ASE (Autonomous System External) or NSSA routes are permitted in a stub area. Each router in a stub area must specify that they are members of the stub area. When specifying that the ABR is a member of the stub area, the ABR will inject a default route into the area.

Routing to external designations from stub areas is based on a default route injected by a stub area’s ABR. A default route is automatically created by the stub area’s ABR. This default route is injected into the stub area to enable other stub routers within the stub area to reach any external routes that are no longer inserted into the stub area.

A stub area can be configured such that the ABR is prevented from sending type 3 summary LSAs into the stub area using the **no-summary** option. In this case, all destinations outside of the stub area are represented by means of a default route.

There are a couple of restrictions on the use of stub areas. Virtual-links cannot be configured through stub areas, and AS boundary routers cannot be placed internal to stub areas.

Use the **area stub** command in OSPF router configuration command mode to configure an area as a stub.

### Stub Area Default Route Cost

A cost value can be set for the default route that is sent into a stub area by an ABR. Configuration of the stub area default route cost is restricted to the ABR attached to this stub area.

Use the **area default-cost** command in OSPF router configuration command mode on the ABR attached to this stub area to configure the stub area default route cost.

## Example

Every router in Areas 1 and 2 are configured for a stub area (Routers 1, 2, and 3 for Area 1 and Routers 5, 6, 7, and 8 for Area 2). Additionally, ABR routers 3, 5, and 6 are also configured with a default-cost to be assigned to the stub area. Router 5 has a lower metric cost when compared to Router 6, so Router 5 will be the preferred router for packets to access the area, with Router 6 employed as a backup in case Router 5 fails. The following example provides the input required to configure the stub topology as displayed in [Figure 32-5](#) on page 32-12:

### Router 1

```
Router1(rw-config)->router ospf 1
Router1(rw-config-ospf-1)->area 1 stub
```

### Router 2

```
Router2(rw-config)->router ospf 1
Router2(rw-config-ospf-1)->area 1 stub
```

### Router 3

```
Router3(rw-config)->router ospf 1
Router3(rw-config-ospf-1)->area 1 stub
Router3(rw-config-ospf-1)->area 1 default-cost 15
```

### Router 5

```
Router5(rw-config)->router ospf 1
Router5(rw-config-ospf-1)->area 2 stub
Router5(rw-config-ospf-1)->area 2 default-cost 15
```

## Router 6

```
Router6(rw-config)->router ospf 1
Router6(rw-config-ospf-1)->area 2 stub
Router6(rw-config-ospf-1)->area 2 default-cost 20
```


## Router 7

```
Router7(rw-config)->router ospf 1
Router7(rw-config-ospf-1)->area 2 stub
```

## Router 8

```
Router8(rw-config)->router ospf 1
Router8(rw-config-ospf-1)->area 2 stub
```

**Figure 32-5 OSPF Stub Area Topology**



## Configuring a Not So Stubby Area (NSSA)

A Not So Stubby Area (NSSA) is a hybrid area using an Autonomous System Border Router (ASBR) to connect two disparate organizations. External routes are advertised as Type 7 LSAs and are converted to Type 5 LSAs before flooding to the backbone by the NSSA's ABR. Also, summary routes are allowed into the NSSA while external routes from other networks are still filtered from insertion into the NSSA.

External routes that are not imported into an NSSA can be represented by a default route. If the router is an ABR and has the highest router ID of all ABRs in the area, and no other ABR in the area is configured to translate always, it will translate Type 7 LSAs into Type 5 LSAs. Configuring the identity of the translator can be used to bias the routing to aggregated destinations. When translator role is set to Always, Type-7 LSAs are always translated regardless of the translator state of other NSSA border routers.

When a translating ABR loses a translator election, it will stop translating, and after a number of seconds (set by the **transstabilityint** option), it will flush any Type 5 LSAs resulting from aggregation. Any Type 5 LSAs resulting from direct translation of Type 7 LSAs will be allowed to age out. An ABR will always originate a default route into any attached NSSAs.

If the **no-summary** option is specified, the ABR does not send type 3 summary LSAs into the NSSA area, therefore all destinations outside of the NSSA area are represented by means of a default route.

Use the **area nssa** command to configure an area as a Not-So-Stubby-Area.

## Example

Routers 2 and 6 are configured as the ABRs between Area 1 and 0, and Router 4 as the ASBR. Router 2 is configured to set Area 1 as an NSSA, and Type 7 routes from the connected domain will be translated to Type 5 routes into the backbone.

ABR Router 2 will only translate Type 7 LSAs; static routes redistributed by router 4. Also, Router 2 will always translate, since it is configured to do so; Router 6 will not, since only one ABR will perform the translation for a given area.

Router 4 will be configured to redistribute static routes.

The following example provides the input required to configure the NSSA topology as displayed in [Figure 32-6](#) on page 32-14:

### Router 6 (ABR)

```
Router 6 (rw) ->configure
Router 6 (rw-config) ->interface vlan 1
Router 6 (rw-config-intf-vlan.0.1) ->ip address 11.1.1.6 255.255.255.252
Router 6 (rw-config-intf-vlan.0.1) ->no shutdown
Router 6 (rw-config-intf-vlan.0.1) ->exit
Router 6 (rw-config) ->interface vlan 2
Router 6 (rw-config-intf-vlan.0.2) ->ip address 23.1.1.6 255.255.255.252
Router 6 (rw-config-intf-vlan.0.2) ->no shutdown
Router 6 (rw-config-intf-vlan.0.2) ->exit
Router 6 (rw-config) ->router ospf 1
Router 6 (rw-config-ospf-1) ->router-id 6.6.6.6
Router 6 (rw-config-ospf-1) ->area 1 nssa
Router 6 (rw-config-ospf-1) ->network 11.1.1.0 0.0.0.3 area 0
Router 6 (rw-config-ospf-1) ->network 23.1.1.0 0.0.0.3 area 1
Router 6 (rw-config-ospf-1) ->exit
```

### Router 2(ABR)

```
Router 2 (rw) ->configure
Router 2 (rw-config) ->interface vlan 1
Router 2 (rw-config-intf-vlan.0.1) ->ip address 11.1.1.2 255.255.255.252
Router 2 (rw-config-intf-vlan.0.1) ->no shutdown
Router 2 (rw-config-intf-vlan.0.1) ->exit
Router 2 (rw-config) ->interface vlan 2
Router 2 (rw-config-intf-vlan.0.2) ->ip address 23.1.1.1 255.255.255.252
Router 2 (rw-config-intf-vlan.0.2) ->no shutdown
Router 2 (rw-config-intf-vlan.0.2) ->exit
Router 2 (rw-config) ->router ospf 1
Router 2 (rw-config-ospf-1) ->router-id 2.2.2.2
Router 2 (rw-config-ospf-1) ->network 11.1.1.0 0.0.0.3 area 0
Router 2 (rw-config-ospf-1) ->network 23.1.1.0 0.0.0.3 area 1
Router 2 (rw-config-ospf-1) ->area 1 nssa
Router 2 (rw-config-ospf-1) ->area 1 nssa transrole always
Router 2 (rw-config-ospf-1) ->area 1 nssa-range 10.2.0.0 255.255.0.0
Router 2 (rw-config-ospf-1) ->exit
```


## Router 4 (ASBR)

```

Router 4 (rw)->configure
Router 4 (rw-config)->interface vlan 2
Router 4 (rw-config-intf-vlan.0.1)->ip address 23.1.1.2 255.255.255.252
Router 4 (rw-config-intf-vlan.0.1)->no shutdown
Router 4 (rw-config-intf-vlan.0.1)->exit
Router 4 (rw-config)->interface vlan 3
Router 4 (rw-config-intf-vlan.0.2)->ip address 30.1.1.1 255.255.255.252
Router 4 (rw-config-intf-vlan.0.2)->no shutdown
Router 4 (rw-config-intf-vlan.0.2)->exit
Router 4 (rw-config)->router ospf 1
Router 4 (rw-config-ospf-1)->router-id 4.4.4.4
Router 4 (rw-config-ospf-1)->network 23.1.1.0 0.0.0.3 area 1
Router 4 (rw-config-ospf-1)->redistribute static metric-type 1
Router 4 (rw-config-ospf-1)->exit

```

**Figure 32-6 OSPF NSSA Topology**



## Configuring Area Virtual-Links

The backbone area 0 cannot be disconnected from any other areas in the AS. Disconnected areas will become unreachable. To establish and maintain backbone connectivity, virtual-links can be configured through non-backbone areas for the purpose of connecting a disconnected area with the backbone through a backbone connected area. The two endpoints of a virtual link are ABRs, both of which belong to the backbone connected area (also referred to as the transit area); one of which belongs to the area disconnected from the backbone. Virtual links cannot be configured through stub areas (see “[Configuring a Stub Area](#)” on page 32-10 for stub area configuration information).

The virtual-link is treated as if it were an unnumbered point-to-point network belonging to the backbone and joining the two ABRs. The cost of a virtual link is not configured. It is auto configured with the cost of the intra-area path between the two ABRs that make up the virtual-link.

Use the **area virtual-link** command in OSPF router configuration command mode, providing the transit area ID and the ABRs IP address, to configure an area virtual-link.

[Figure 32-7](#) on page 32-15 displays a typical virtual-link topology. Area 3 does not share an ABR with the backbone area, and is therefore disconnected from the backbone. Area 3 shares an ABR

(router 2) with area 1. Area 1 has a second ABR (router 1) that it shares with the backbone. Area 1 is the transit area because it contains an ABR that it shares with the disconnected area and a second ABR that it shares with the backbone. By configuring an area virtual-link between router 2 and router 1, Area 3 will gain connectivity with the backbone and be able to learn routes for this AS.

## Example

The following example presents the configuration required to configure the virtual-link displayed in [Figure 32-7](#) on page 32-15:


### Router 1

```
Router 1 (rw-config)->router ospf 1
Router 1 (rw-config-ospf-1)->area 0.0.0.1 virtual-link 2.2.2.2
Router 1 (rw-config-ospf-1)->exit
Router 1 (rw-config)->
```

### Router 2

```
Router 2 (rw-config)->router ospf 2
Router 2 (rw-config-ospf-2)->area 0.0.0.1 virtual-link 1.1.1.1
Router 2 (rw-config-ospf-2)->exit
Router 2 (rw-config)->
```

**Figure 32-7 Virtual Link Topology**



## Configuring Area Virtual-Link Authentication

An area virtual-link can be configured for either simple or MD5 authentication.

Use the **area virtual-link authentication-key** command in OSPF router configuration command mode to configure simple authentication on this area virtual-link.

Use the **area virtual-link message-digest-key** command in OSPF router configuration command mode to configure MD5 authentication on this area virtual-link.

### Configuring Area Virtual-Link Timers

The following timers can be configured for an area virtual-link:

- Dead-interval using the **area virtual-link dead-interval** command
- Hello-interval using the **area virtual-link hello-interval** command
- Retransmit-interval using the **area virtual-link retransmit-interval** command
- Transmit-delay using the **area virtual-link transmit-delay** command

See “[Configuring OSPF Timers](#)” on page 32-20 for an OSPF timers discussion.



**Note:** RFC 2328 specifies that the retransmit-interval should be greater than the expected round-trip delay between the two routers. This may be hard to estimate for a virtual link; it is better to err on the side of making it too large.

## Configuring Route Redistribution

Redistribution permits the importation of other routing protocols into OSPF such as RIP, as well as static and directly connected routes. Alternately, you can specify a route-map for redistribution into OSPF. Be aware that if the referenced route map has not yet been configured, then an empty route map is created with the specified name. See “[Configuring a Not So Stubby Area \(NSSA\)](#)” on page 32-12 for an example of redistribution of static routes by an ASBR in an NSSA context.

Use the **redistribute** command in OSPF router configuration command mode to permit the redistributions of OSPF, RIP, static, or connected routes by this router.

## Filtering Routes from the OSPF Route Table

Routes can be filtered from the OSPF route table by creating an OSPF filter route route-map and assigning it to the distribute-list for this OSPF router.

For example, the 10.1.1.0/24 network is advertised via OSPF from Area 0. However, private networks exist in 10.0.0.0/8. Various routers will learn the 10.1.1.0/24 route via OSPF, but they should not route packets to the 10.0.0.0/8 network. The solution is to not allow the 10.1.1.0/24 route to be installed in the forwarding tables by filtering it from the routing table with a route-map based upon its network address.

Use the **route-map filter** command as described in the “Route-Map Manager” section of the *Extreme Networks 7100-Series CLI Reference* to create an OSPF filter route route-map.

Use the **distribute-list route-map in** command to assign the filter route-map to the OSPF distribute-list.

## Configuring Passive Interfaces

Passive interfaces explicitly allows the network to be advertised, but prevents it from forming neighbor relationships on that interface. Passive interfaces are included in the OSPF route table. They do not send or receive hello packets. OSPF adjacencies can not be formed on a passive interface.

An option exists to default all interfaces to passive mode.

Use the **passive-interface** command in router configuration command mode to configure an interface as passive.

## Graceful Restart

OSPF graceful restart, sometimes referred to as non-stop forwarding, provides for an OSPF router to remain on the forwarding path during a restart of its OSPF software. Graceful-restart has three elements to its configuration: enabling, helper router, and restart interval.

Enabling graceful restart instructs the firmware to perform a graceful restart, rather than a standard OSPF restart. Restart is only initiated by a fail-over. Grace LSAs are sent when OSPF is restarted on another module. Whether the failover is intentional or not, the failed router protocol is restarted on another module, and upon startup, OSPF sends grace LSAs out to its neighbors using existing link aggregation groups. Use the **graceful-restart enable** command to enable the graceful restart ability on this router.

The helper relationship with the restarting router is on a per network segment basis. The helper monitors the network for topology changes. If no changes occur, the helper router continues to advertise its LSAs as though no restart was occurring. If the restarting router was the designated router, the helper continues to treat it as such. If a topology change does occur, graceful restart is terminated on the restarting router and a standard restart occurs. Helper mode can be disabled on a restarting router neighbor using the **ip ospf helper-disable** command in interface command mode. If the restarting router receives an LSA indicating a disabled helper, the graceful restart terminates and a standard restart occurs.

A restart interval provides for a maximum time in seconds after which the graceful restart will terminate should it not complete or terminate for other reasons within the interval. Use the **graceful-restart restart-interval** command to change the restart interval setting.

View the router OSPF section of the **show running-config** display to verify any non-default graceful restart settings.

## Graceful Restart and High Availability

The 7100-Series module supports single router high availability failover using the following components:

- OSPF graceful restart
- Non-stop router frame forwarding on each module
- Single router configuration
- Router protocol process failover to another module
- Link Aggregate Group (LAG) connectivity to neighboring routers

In a stable network, the route and rule information is fairly constant. If the router protocol process was to suddenly fail, forwarding information current at the time of the failure in all probability is usable for the short time after the failure until recovery occurs. During this recovery period, existing connections (that were not directly using the failed module) remain in effect. New connections continue to be installed using the last known “good” forwarding information. The router protocol process that failed is dynamically restarted. The user does not configure where the router process is running. The router forwarding process remains active on every module. The protocol process exchanges protocol and maintains state that it distributes to the other modules and does not have to run on any specific module. One exception to this rule is that the module must have 256M of memory to be router protocol process eligible.

Upon failure of a module running the router protocol process, the protocol process is started on a recovery module. One of the first messages it sends to its OSPF neighbors is a grace LSA. High availability failover will successfully occur if the following is true:

- The router is enabled for graceful restart
- The neighbors are enabled to participate as graceful restart helper



- The OSPF dead interval is configured for a sufficient period such that the grace LSA is received by its neighbors before the configured OSPF dead interval expires
- And each neighbor is a member of a LAG common to the failed router, allowing the neighbor to remain up

**Figure 32-8 Physical and Logical Single Router HA Failover Configuration**




Figure 32-8 depicts the physical and logical configurations of the single router high availability failover mechanism. The neighbor to router lines display direct neighbor connections to the router enabled for OSPF graceful restart and members of LAGs common to the failing router. The server to router lines display VLAN connections common to both the failing and recovery routers.

## Configuring Interface Cost

Each interface has an outbound cost associated with it. The lower the cost, the more likely the interface will be used to forward data traffic. Should several equal-cost routes to a destination exist, traffic is distributed equally among them.

The formula for calculating the OSPF interface cost metric is the reference bandwidth divided by the interface bandwidth. By default the reference bandwidth is set to 100 Mbps. For 10 Mbps links, the resulting cost is 10. For 100, 1000, or 10000 Mbps links, the resulting cost is 1. The reference bandwidth can be modified using the **auto-cost reference-bandwidth** command in OSPF configuration mode. The ability to re-center the reference bandwidth to a higher value, allows for OSPF interface costs to default to a value greater than 1 for 100, 1000, or 10000 Mbps links and greater than 10 for 10 Mbps links.

It is recommended that the auto cost reference bandwidth be the same value for all OSPF routers in the domain.

Use the **ip ospf cost** command in interface configuration command mode to statically specify the outbound cost of this interface. A statically configured OSPF interface cost overrides all other interface cost methods.

For logical interfaces containing multiple physical interfaces, such as a LAG, the aggregate interface speed is not readily available. A tracked object configured with the ports belonging to the logical interface can return the physical interface speed of each physical port specified in the tracked object. OSPF will sum the returned interface speeds and use that aggregate value when calculating OSPF interface cost. Because the tracked object will report when a physical interface is up or down, OSPF will dynamically adjust the aggregate speed when an interface becomes active or goes down and adjust the OSPF interface cost accordingly. This method should be used in LAG and ECMP logical interface contexts.



**Note:** The speed used in the cost calculation is sum of all ports capabilities in the tracked object. Setting the speed manually will not change the tracked interface speed. A 1GB capable port has a 1 GB speed regardless of the manual speed setting. The same holds true for ports that auto-negotiate to a lower speed. The expectation is that both sides of the link are using the same ports and SFP connectors and should result in the same speed.

Use the **ip ospf cost track** command in interface configuration mode to calculate the OSPF interface cost based upon summing physical interface speeds that belong to a logical interface.

When adding an additional physical port to a logical interface that uses the interface summation method to determine OSPF interface cost, you must also add the physical port to the associated tracked object.

See “[Tracked Object Manager Configuration](#)” on page 12-1 for tracked object configuration details.

## Configuring OSPF with Authentication at the Interface

Authentication helps ensure that routing information is processed only from trusted routers. This section describes OSPF authentication at the interface level. Two authentication schemes can be used, simple using the **ip ospf authentication-key** command or MD5 using the **ip ospf message digest key md5** command, but a single scheme must be configured for each network. The use of different schemes enables some interfaces to use much stricter authentication than others. When you wish to bar routers from exchanging OSPF packets, use simple authentication. The interfaces that the packets will be sent on still must be trusted because the authentication key will be placed in the packets and are visible to anyone on the network. An adjacency with another router will not occur unless the simple authentication is configured the same on both ends of the interface.

If you do not trust other routers on your network, use MD5 authentication. The system works by using shared secret keys. Because keys are used to sign the packets with an MD5 checksum through a one-way hash function, they cannot be forged or tampered with. Also, because the keys are not included in the packet, snooping the key is impossible. Network users can still snoop the contents of packets, though, because packets are not encrypted.

7100-Series device MD5 authentication is compliant with OSPF RFC 2328. This specification uses the MD5 algorithm and an authentication key of up to 16 characters.

## Configuring Bidirectional Forwarding Detection (BFD) on Interfaces

BFD is used to detect a communications failure with an OSPF forwarding plane next-hop. BFD detects failures in under one second. BFD augments the OSPF Hello mechanism. The OSPF Hello interval defaults to 10 seconds. With high speed data rates, a failure requiring multiple seconds to detect can result in significant data loss. The OSPF implementation of the BFD protocol uses the following non-configurable parameters:

**Transmit Interval** – The period of time between the transmission of BFD control packets, set for 100ms.

**Receive Interval** – The period of time between received BFD control packets, set for 100ms.

**Detection Multiplier** – The Number of consecutive control packets that can be missed before the BFD session transitions to down, set to 3.

Use the **bfd all-intfs-on** command in OSPF router configuration mode to enable BFD on all OSPF interfaces.

## Configuring OSPF Timers

There are five OSPF timers:

- Hello-Interval
- Dead-Interval
- Retransmit-Interval
- Transmit-Delay
- SPF-Delay

To ensure efficient adjacency between OSPF neighbors, the 7100-Series device provides hello-interval and dead-interval commands. The hello interval is the period between transmissions of hello packet advertisements. The dead interval is the period that can elapse without receiving a router's hello packets before its neighbors will declare it down.

Use the **ip ospf hello-interval** command in interface configuration command mode to configure the period between transmissions of hello packet advertisements.

Use the **ip ospf dead-interval** in interface configuration command mode to configure the period between receiving hello packets before the associated neighbor is declared down.

In order to ensure that flooding is reliable, LSAs are retransmitted until they are acknowledged. The period between retransmissions is the retransmit-interval. If this interval is set too low for an interface, needless retransmissions will take place. If the value is set too high, the speed of the flooding, during the period of lost packets, may be affected.

Use the **ip ospf retransmit-interval** command in interface configuration command mode to configure the retransmit-interval.

The transmit-delay is an estimation of the number of seconds it takes to transmit a link state update packet over this interface. This value should take into account transmission and propagation delays.

Use the **ip ospf transmit-delay** command in interface configuration command mode to configure the transmit-delay.

The SPF-delay is the amount of time that transpires between the receipt of an OSPF update and the SPF calculation.

Use the **timers spf** command in OSPF router configuration command mode to specify the amount of time between receiving an OSPF update and an SPF calculation occurring.

The OSPF timers can also be configured for an area virtual-link. See "[Configuring Area Virtual-Links](#)" on page 32-14.

## Configuring the PE-CE Protocol

The PE-CE protocol allows a service provider offering Virtual Private Network (VPN) services to their customers to peer Customer Edge (CE) routers with their Provider Edge (PE) routers. RFC 4577 defines how the PE-CE protocol is implemented using the OSPF routing protocol.

When the PE router becomes a routing peer of the CE router, the PE router learns the routes that lead to the CE's site and can redistribute those routes to other PE routers that attach to the same VPN.

Enabling PE-CE enables the following functionality:

- DN Bit
- Sham link
- Domain tag
- Domain ID

Use the **enable-pe-ce** command in OSPF configuration mode to enable the PE-CE protocol on the router.

## The OSPF VRF Domain Tag

The configuration and inclusion of the OSPF VRF domain tag is required for PE-CE protocol enabled systems to be backward compatible with systems that do not set the PE-CE protocol DN bit in type 5 LSAs. When a prefix is received from a BGP speaker and redistributed into the PE-CE protocol enabled OSPF instance, the OSPF process for the VRF is given a domain tag. In the event that the customer site attempts to re-advertise the prefix to another PE using the same domain tag, the domain tag will be matched and the prefix will not be accepted by the second PE for redistribution into BGP. Setting the same domain tag for all backbone PE routers on the same VPN prevents routing loops.

Setting the OSPF VRF domain tag is optional when the PE-CE protocol is enabled for all PE backbone routers for a given VRF. If legacy PE routers that do not support the PE-CE protocol are present in the VRF backbone, set the domain tag for this router to agree with the domain tag of the legacy router.

The PE-CE protocol must be enabled using the **enable-pe-ce** command to set the OSPF VRF domain tag.

Use the **domain-tag** command in OSPF configuration mode to set the domain-tag for this PE router.

## The OSPF VRF Domain ID

If the OSPF instances of an OSPF domain are given one or more domain IDs, OSPF can determine whether an OSPF-originated VPN-IPv4 route belongs to the same domain as a given OSPF instance and whether the route should be redistributed to that OSPF instance as an inter-area route or as an OSPF AS-external route.

If two OSPF instances with a domain ID configured are in the same OSPF domain, the PE-CE protocol requires that the primary domain ID of the second instance must be one of the domain IDs of the first instance (either primary or secondary). If two OSPF instances with a domain ID configured are not in the same OSPF domain, the primary domain ID of each instance must not be configured as a domain ID of the other OSPF instance.

The PE-CE protocol must be enabled using the **enable-pe-ce** command to set the OSPF VRF domain ID.

Use the **domain-id** command in OSPF configuration mode to set the domain ID for this PE router.

## Redistribute BGP into OSPF

OSPF supports the redistribution of BGP discovered routes into OSPF. A **global** option is available for the redistribution of BGP prefixes from the global router. When using the **global** option, VPN4 address prefixes are appropriately translated.

Use the **redistribute bgp global** command in OSPF configuration mode to redistribute BGP learned routes to other PE routers in this VPN.

## OSPF Sham Link

If a VRF contains both an OSPF-distributed route and a VPN-IPv4 route for the same IPv4 prefix, then the backdoor OSPF-distributed route is preferred over the VPN backbone route, unless the next hop interface for an installed (OSPF distributed) route is the sham link, in which case, the VPN backbone VPN-IPv4 route is used.

If it is desired to have OSPF prefer the routes through the VPN backbone over the routes through the OSPF backdoor link, then the routes through the backbone must appear to be intra-area routes. The sham link provides this appearance of an intra-area link connecting the two PE routers.

Use the **area sham-link** command in OSPF configuration mode, specifying both the source and destination link addresses to configure an OSPF sham link between two VPN PE routers.

## Configuring OSPF

This section provides details for the configuration of OSPF on 7100-Series platforms.

### Default Settings

[Table 32-1](#) lists OSPF parameters and their default values.

**Table 32-1 Default OSPF Parameters**

Parameter	Description	Default Value
router ID	Provides for the identification of one router to another and helps establish adjacencies among OSPF routers.	highest IP address of configured routing interfaces
interface cost	An outbound interface value used in determining which routing interface should forward when more than one routing interface is available.	10
interface priority	A value placed on the interface that helps in determining which router will be elected designated router.	1
interface network type	Specifies the type of network an interface is connecting to.	broadcast
LSA Thresholds	Specifies : <ul style="list-style-type: none"> <li>The number of LSA updates that force a full routing calculation</li> <li>The number of LSA updates that interrupt and restart a full routing calculation</li> <li>The number of LSA inter-area/external updates that force a full routing calculation</li> <li>the number of intra updates that force a full routing calculation</li> </ul>	4294967295 Update starts 4294967295 Update restarts 50 Inter-area/external updates 0 Intra updates
LSA Pause Frequency	Specifies the number of units SPF calculation runs before pausing.	10000

**Table 32-1 Default OSPF Parameters (continued)**

Parameter	Description	Default Value
SPF delay timer	Specifies the amount of time between receiving an OSPF update and the start of an SPF calculation.	5 seconds
retransmit interval	A timer that determines the retransmission of LSAs in order to ensure reliable flooding.	5 seconds
transmit delay	Specifies the number of seconds it takes to transmit a link state update packet over this interface.	1 second
hello interval	The period between transmissions of hello packet advertisements.	10 seconds for broadcast and point-to-point networks; 30 seconds for non-broadcast and point-to-multipoint networks
dead interval	The period that can elapse without receiving a router's hello packets before its neighbors will declare it down.	40 seconds
distance	Specifies the administrative distance for OSPF routes. The available protocol with the lowest administrative distance is chosen for this route.	connected = 0 static = 1 OSPF = 110 RIP = 120
graceful-restart	Provides for an OSPF router to remain on the forwarding path during a restart of its OSPF software.	disabled
graceful-restart restart interval	Specifies the maximum time in seconds after which the graceful restart will terminate should it not complete or terminate for other reasons within the interval.	120 seconds
PE-CE Protocol	Enables the Customer Edge (CE) router as a Provider Edge (PE) router peers.	disabled

[Procedure 32-1](#) describes how to configure basic OSPF parameters. All commands in this procedure are entered in OSPF router configuration command mode, except where indicated.

### Procedure 32-1 Configuring Basic OSPF Parameters

Step	Task	Command(s)
1.	Configure an IP address for all routing interfaces in the AS. <ul style="list-style-type: none"> <li>• <b>primary</b> - (Optional) Specifies that the configured IP address is a primary address.</li> <li>• <b>secondary</b> - (Optional) Specifies that the configured IP address is a secondary address.</li> </ul>	<b>ip address</b> { <i>ip-address</i>   <i>ip-address/prefixLength</i> } <i>ip-mask</i> [ <b>primary</b>   <b>secondary</b> ]
2.	Create an OSPF routing instance.	<b>router ospf</b> <i>process-id</i>

**Procedure 32-1 Configuring Basic OSPF Parameters (continued)**

Step	Task	Command(s)
3.	Configure the network addresses, masks, and areas for each subnet on this AS. <ul style="list-style-type: none"> <li>• <b>area</b> - Specifies the <i>area-id</i> to be associated with the OSPF address range. Valid values are decimal values between <b>0 - 4294967295</b> or an IP address. A subnet address can be specified as the <i>area-id</i> to associate areas with IP subnets.</li> </ul>	<b>network</b> <i>ip-address wildcard-mask area area-id</i>

[Procedure 32-2](#) describes how to configure basic OSPF parameters.

**Procedure 32-2 Configuring OSPF General Optional Parameters**

Task	Command(s)
Optionally, change the OSPF router ID for this device.	<b>router-id</b> <i>ip-address</i>
Optionally, enable the OSPF PE-CE protocol.	<b>enable-pe-ce</b>
Optionally, configure the OSPF VRF domain tag.	<b>domain-tag</b> <i>tag</i>
Optionally, configure the OSPF VRF domain ID.	<b>domain-id [secondary] type type value value</b>
Optionally, configure the OSPF router neighbors for this router.	<b>neighbor</b> <i>ip-address [priority priority]</i>
Optionally, change the SPF LSA thresholds for this router.	<b>spf lsa-thresholds</b> <i>num-start num-restart num-intra-full num-ia-ext-full</i>
Optionally, change the SPF pause frequency to specify the number of units SPF calculation runs before pausing.	<b>spf pause-frequency</b> <i>units</i>
Optionally, change the delay, in milliseconds, between the receipt of an update and the beginning of the SPF execution.	<b>timers spf</b> <i>spf-delay</i>
Optionally, enable BFD on all OSPF interfaces.	<b>bfd all-intfs-on</b>
Optionally, change the administrative distance for OSPF routes.	<b>distance [ospf {external   intra-area}] weight</b>
Optionally, define the range of addresses used by this Area Border Router (ABR) when communicating routes to other areas.	<b>area area-id range</b> <i>ip-address ip-mask [not-advertised]</i>
Optionally, configure an area as a stub area.	<b>area area-id stub [no-summary]</b>
Optionally, set the cost for the default route that is sent into a stub area by an ABR.	<b>area area-id default-cost</b> <i>cost</i>
Optionally, configure an area as a not so stubby area.	<b>area {area-id   ip-address} nssa [no-summary] [transstabilityint seconds] [transrole always]</b>
Optionally, configure an Autonomous System Border Router (ASBR) to summarize Type 7 to Type 5 routes matching the specified address and mask.	<b>area {area-id   ip-address} nssa-range</b> <i>ip-address mask</i>
Optionally, configure an OSPF sham link between two PE routers.	<b>area area-id sham-link</b> <i>source-ip-address destination-ip-address</i>

**Procedure 32-2 Configuring OSPF General Optional Parameters (continued)**

<b>Task</b>	<b>Command(s)</b>
Optionally, configure an OSPF sham link authentication key password.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>authentication-key</b> <i>password</i>
Optionally, modify the OSPF sham link dead interval.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>dead-interval</b> <i>seconds</i>
Optionally, modify the OSPF sham link hello interval.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>hello-interval</b> <i>seconds</i>
Optionally, configure the OSPF sham link keychain.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>keychain</b> <i>name</i>
Optionally, specify an OSPF sham link message digest key and MD5 authentication key.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>message-digest-key</b> <i>digest-key</i> <b>md5</b> <i>auth-key</i>
Optionally, modify the OSPF sham link retransmit interval.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>retransmit-interval</b> <i>seconds</i>
Optionally, modify the OSPF sham link transmit delay period.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>transmit-delay</b> <i>seconds</i>
Optionally, modify the OSPF sham link cost.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>cost</b> <i>cost</i>
Optionally, configure an OSPF virtual-link, which represents a logical connection between the backbone and a non-backbone OSPF area.	<b>area</b> <i>area-id</i> <b>virtual-link</b> <i>ip-address</i> <b>area</b> <i>area-id</i> <b>virtual-link</b> <i>ip-address</i> <b>authentication-key</b> <i>key</i> <b>area</b> <i>area-id</i> <b>virtual-link</b> <i>ip-address</i> <b>dead-interval</b> <i>seconds</i> <b>area</b> <i>area-id</i> <b>virtual-link</b> <i>ip-address</i> <b>hello-interval</b> <i>seconds</i> <b>area</b> <i>area-id</i> <b>virtual-link</b> <i>ip-address</i> <b>message-digest-key</b> <i>digest-key</i> <b>md5</b> <i>format</i> <i>line</i> <i>auth-key</i> <b>area</b> <i>area-id</i> <b>virtual-link</b> <i>ip-address</i> <b>retransmit-interval</b> <i>seconds</i> <b>area</b> <i>area-id</i> <b>virtual-link</b> <i>ip-address</i> <b>transmit-delay</b> <i>seconds</i>
Optionally, change the bandwidth reference setting used for calculating interface cost for this OSPF instance.	<b>auto-cost</b> <b>reference-bandwidth</b> <i>bandwidth-multiplier</i>
Optionally, enable passive OSPF on the specified interface.	<b>passive-interface</b> { <i>vlan-id</i>   <i>interface-name</i>   <b>default</b> }
Optionally, allow routing information discovered through non-OSPF protocols to be distributed in OSPF update messages.	<b>redistribute</b> { <b>rip</b>   <b>static</b>   <b>connected</b> } [ <b>route-map</b> <i>id-number</i> ] [ <b>metric</b> <i>metric value</i> ] [ <b>metric-type</b> <i>type-value</i> ] [ <b>tag</b> <i>tag</i> ]
Optionally, assign an OSPF route filter route-map to the OSPF distribute-list.	<b>distribute-list</b> <b>route-map</b> <i>name</i> <b>in</b>
Optionally, enable the graceful-restart feature on this router.	<b>graceful-restart</b> <b>enable</b>
Optionally, change the graceful-restart restart interval for this router.	<b>graceful-restart</b> <b>restart-interval</b> <i>interval</i>



**Procedure 32-2 Configuring OSPF General Optional Parameters (continued)**

Task	Command(s)
Optionally, in system command mode, reset the specified OSPF process ID or the OSPF process.	<b>clear ip ospf process</b> [ <i>process-id</i> ]
Optionally, in global configuration command mode, enable OSPF protocol debugging output for the specified subsystem.	<b>debug ip ospf</b> { <i>subsystem</i> }
Optionally, enable this OSPF router for RFC 1583 compatibility.	<b>rfc1583compatible</b>

[Procedure 32-3](#) describes how to configure optional OSPF interface parameters. All commands in this procedure are entered in interface configuration command mode.

**Procedure 32-3 Configuring OSPF Optional Interface Parameters**

Step	Task	Command(s)
1.	Optionally, change the cost of sending an OSPF packet on this router interface. This setting overrides all other interface cost methods.	<b>ip ospf cost</b> <i>cost</i>
2.	Optionally, sum the interface speeds contained in the specified tracked object when calculating the OSPF interface cost.	<b>ip ospf cost track</b> <i>trackobject-name</i>
3.	Optionally, change the OSPF priority value for this router interface.	<b>ip ospf priority</b> <i>number</i>
4.	Optionally, change the OSPF poll-interval value for this non-broadcast neighbor.	<b>ip ospf poll-interval</b> <i>seconds</i>
5.	Optionally, change the amount of time between retransmissions of LSAs for adjacencies that belong to this interface.	<b>ip ospf retransmit-interval</b> <i>seconds</i>
6.	Optionally, change the amount of time required to transmit a link state update packet on this interface.	<b>ip ospf transmit-delay</b> <i>seconds</i>
7.	Optionally, enable the ignore MTU advertisement feature for the neighbor of this interface.	<b>ip ospf ignore-mtu</b>
8.	Optionally, change the number of seconds this router must wait before sending a hello packet to neighbor routers on this interface.	<b>ip ospf hello-interval</b> <i>seconds</i>
9.	Optionally, change the number of seconds this router must wait to receive a hello packet from its neighbor before determining that the neighbor is out of service.	<b>ip ospf dead-interval</b> { <i>seconds</i>   <b>minimal hello-multiplier</b> <i>number</i> }
10.	Optionally, assign a password on this interface to be used by neighboring routers using OSPF's simple password authentication.	<b>ip ospf authentication-key</b> <i>password</i>
11.	Optionally, enable OSPF MD5 authentication on this interface.	<b>ip ospf message-digest-key</b> <i>keyid md5 key</i>
12.	Optionally, disable the graceful restart helper feature on this interface.	<b>ip ospf helper-disable</b>

**Procedure 32-3 Configuring OSPF Optional Interface Parameters (continued)**

Step	Task	Command(s)
13.	Optionally, specify the network type that this interface is connected to.	<b>ip ospf network {non-broadcast   broadcast   point-to-point   point-to-multipoint}</b>

[Table 32-2](#) describes how to display OSPF configuration and statistics.

**Table 32-2 Displaying OSPF Configuration and Statistics**

Task	Command(s)
Displaying OSPF configuration.	<b>show ip ospf</b>
Displaying OSPF link state database information.	<b>show ip ospf database</b> [ <i>link-state-id</i> ]
Displaying information about OSPF internal entries to area border routers and autonomous system boundary routers.	<b>show ip ospf border-routers</b>
Displaying OSPF interface configuration information.	<b>show ip ospf interface</b> [ <i>vlan vlan-id</i> ]
Displaying OSPF neighbor information.	<b>show ip ospf neighbor</b> [ <i>detail</i> ] [ <i>ip-address</i> ] [ <i>vlan vlan-id</i> ]
Displaying OSPFv3 sham link information.	<b>show ipv6 ospf sham-link</b>
Displaying OSPF virtual-links configuration information.	<b>show ip ospf virtual-links</b>



## Open Shortest Path First Version 3 (OSPFv3) Configuration

This chapter provides the following information about configuring and monitoring OSPFv3 on Extreme Networks 7100-Series devices:

For information about...	Refer to page...
<a href="#">Using the OSPFv3 Protocol in Your Network</a>	33-1
<a href="#">Implementing OSPFv3</a>	33-5
<a href="#">OSPFv3 Configuration Overview</a>	33-5
<a href="#">OSPFv3 Configuration Details</a>	33-25



**Note:** OSPFv3 requires an advanced router license. See “[License Overview](#)” on page 9-10 for 7100-Series advanced router license details.

### Using the OSPFv3 Protocol in Your Network

Open Shortest Path First Version 3 (OSPFv3) is the OSPF routing protocol for IPv6. OSPFv3 is considered a TCP/IP internet routing Interior Gateway Protocol (IGP). OSPFv3 distributes routing information between routers belonging to a single Autonomous System (AS). The OSPF protocol is based on link-state or SPF technology. The advantages associated with a link-state routing protocol are:

- Rapid convergence
- Reduced routing update traffic over traditional distance-vector protocols

This OSPFv3 implementation supports RFC 2740 *OSPF for IPv6*.

The OSPFv3 protocol is designed expressly for the TCP/IP internet environment. OSPFv3 utilizes IP multicast when sending and receiving routing updates. Routing updates are optionally authenticated using IPsec for OSPFv3.

OSPFv3 routes IP packets based solely on the destination IP address found in the IP packet header. IP packets are not encapsulated in any further protocol headers as they transit the AS. OSPFv3 is a dynamic routing protocol in that it quickly detects topological changes in the AS, such as router interface failures, and calculates new loop-free routes after a period of convergence. This period of convergence is short and involves a minimum of routing traffic. In a link-state routing protocol, each router maintains a database describing the AS's topology. This database is referred to as the link-state database. Each participating router has an identical database. Each individual database entry is a particular router's local state made up of such information as the router's usable interfaces and reachable neighbors. The router distributes its local state throughout the AS by flooding.

Each network that has at least two attached routers has a designated router. The designated router generates an LSA for the network and has other special responsibilities in the running of the protocol, enabling a reduction in the number of adjacencies required on a network. This in turn reduces the amount of routing protocol traffic and the size of the link-state database.

All routers run the exact same algorithm, in parallel. From the link-state database, each router constructs a tree of shortest paths with itself as root. This shortest-path tree provides the route to each destination in the AS. Externally derived routing information appears on the tree as leaves. When several equal-cost routes to a destination exist, traffic is distributed equally among them. The cost of a route is described by a single dimensionless metric.

OSPF allows sets of networks to be grouped together. Such a grouping is called an area. The topology of an area is hidden from the rest of the AS. This information hiding enables a significant reduction in routing traffic. Also, routing within the area is determined only by the area's own topology, lending the area protection against bad routing data. An area is a generalization of an IP subnetted network. OSPF enables the flexible configuration of IP subnets. Each route distributed by OSPF has a destination. Two different subnets of the same IP network number may have different masks providing a different range of addresses for that subnet. This is commonly referred to as Variable Length Subnet Masking (VLSM). A packet is routed to the longest or most specific match. Host routes are considered to be subnets whose masks are "all ones" (0xffffffff).

If IPsec for OSPFv3 is enabled on the interface, OSPFv3 protocol exchanges are authenticated. This means that only trusted routers can participate in the AS's routing. The 7100-Series platform supports IPsec for OSPFv3. See ["IPsec for OSPFv3"](#) on page 33-4 for a listing of supported authentication and encapsulation algorithms.

Route redistribution is supported for RIP, connected, BGP, and static routes.

The Bidirectional Forwarding Detection (BFD) protocol providing sub-second failure detection on OSPF forwarding interfaces is enabled by default on all OSPF interfaces.

An OSPF Customer Edge (CE) router can be configured as a peer to a Provider Edge (PE) router by enabling the PE-CE protocol on the PE-CE associated routers.

OSPFv3 is similar to OSPFv2 in its usage of the SPF algorithm, flooding, Designated Router (DR) election, timers, metrics, concept of a link-state database, intra/inter area and AS external routes and virtual-links. OSPFv3 differs with OSPFv2 in many respects, as outlined in ["OSPFv3 and OSPFv2 Differences"](#) on page 33-2.

OSPFv3 is not backward compatible with OSPFv2. If you need to route both IPv4 and IPv6 using OSPF, enable both OSPFv3 and OSPFv2 on the device.

## OSPFv3 and OSPFv2 Differences

OSPFv3 differs from OSPFv2 in a number of respects:

- OSPFv3 processing is per link. OSPFv3 is link rather than subnet centric. An OSPFv3 interface connects to a link, not a subnet. This change has both functional and efficiency advantages. Multiple IP subnets can be assigned to a single link. Two nodes can talk directly over a single link. In OSPFv3, terms such as network and subnet should generally be replaced with the term link in order to understand OSPFv3 processing.
- With the exception of the new link LSA, OSPFv3 LSAs do not carry IPv6 addresses. The removal of addressing from the LSAs has scaling advantages. Router and network LSAs now only contain topology information necessary for SPF processing and no longer contain network addresses. OSPFv3 LSAs do maintain 32-bit RIDs and LSA IDs. Because OSPFv3 IDs are still expressed in dotted-quad notation, OSPFv3 networks can be easily overlaid on an OSPFv2 network.

- All OSPFv3 neighbors are identified by the neighbor router ID. The OSPFv2 behavior of identifying neighbors on broadcast and Non-Broadcast-Multi-Access (NBMA) links by their interface address has been removed.
- OSPFv3 uses link-local addresses, which begin with FF80::/10, as source and next-hop addresses. Link-local addresses are for use on a single link for purposes including neighbor discovery and auto-configuration. On all interfaces, except virtual, OSPF packets are sent using the interface's link-local unicast address as the source. Because all OSPF traffic transits the network on a link basis, IPv6 does not forward (route) IPv6 datagrams having link-local source addresses.
- OSPF specific authentication has been removed and replaced by optionally configuring IPsec for OSPFv3 as defined in RFC 4552. If IPsec for OSPFv3 is not enabled on the interface, OSPFv3 authentication does not take place for OSPF packets.
- RFC 1583 compatibility does not apply to OSPFv3.
- To take advantage of IPv6's link-local scope, OSPFv3 adds a link-local flooding scope to the domain and area flooding scopes present in OSPFv2. The link LSA, which has link-local flooding scope and can not be flooded beyond any attached router, has been added for neighbors on a single link.
- Two new LSAs have been introduced: the link LSA and the intra-area LSA. Point-to-point links are supported in order to enable operation over tunnels. OSPFv3 views IPv6-over-IPv4 tunnels as a point-to-point interface with a link-local address and possibly a global unicast address. OSPFv3 uses the reported MTU for tunnel interfaces.

The prefix advertisement for OSPFv3 is now in the new intra-area prefix LSA. When information is only relevant to the connected neighbor, OSPFv3 puts it in the link LSA, not in the router or network LSA, in both cases avoiding flooding information beyond the relevant information scope.

[Table 33-1](#) details the supported LSA types by LS ID and name for OSPFv3 and OSPF v2.

**Table 33-1 OSPFv3 and OSPFv2 LSA Cross-Reference**

OSPFv3 LSAs	OSPFv2 LSAs
0x2001 – Router LSA	1 – Router LSA
0x2002 – Network LSA	2 – Network LSA
0x2003 – Inter-Area Prefix LSA	3 – Network Summary LSA
0x2004 – Inter-Area Router LSA	4 – ASBR Summary LSA
0x2005 – AS-External LSA	5 – AS-External LSA
0x2006 – Group Membership LSA	6 – Group Membership LSA
0x2007 – Type-7 LSA	7 – NSSA External LSA
0x2008 – Link LSA	No Corresponding LSA for OSPFv2
0x2009 – Intra-Area Prefix LSA	No Corresponding LSA for OSPFv2

- Unlike for OSPFv2, the router and network LSAs for OSPFv3 do not advertise prefixes. In OSPFv3, the router and network LSAs only represent the router's node information for SPF and are only flooded if information relevant to the SPF algorithm changes. This behavior avoids the flooding of prefix changes that are not relevant to SPF.

Inter-area prefix, inter-area router, and type-7 LSAs have the same function as their OSPFv2 counterparts listed in [Table 33-1](#).

- OSPFv3 specifies the processing of unsupported LSAs. Unsupported LSAs are maintained in the database and flooded according to scope. In OSPFv3, routers with 100 or more interfaces generate more than one router LSA. A new link LSA has been created. Addresses in LSAs are specified as [prefix, prefix length].
- OSPFv2 supports multiple OSPF processes on a device. The current OSPFv3 implementation supports a single OSPFv3 process.
- OSPFv3 supports multiple OSPFv3 instances on an interface. Multiple OSPFv3 instances provide for the sharing of an interface when more than one physical network segment needs access to an interface. It also provides for the configuration of multiple areas on a single interface. The multiple OSPF instances feature is not supported by OSPFv2.
- The IPv6 all SPF routers multicast address is FF02::5; the all DRouters multicast address is FF02::6. Both have link-local scope.

## OSPFv3 and OSPFv2 Similarities

- OSPFv3 uses the same 5 message types, with the same message numbering, as OSPFv2:
  - 1 – Hello
  - 2 – Database Description
  - 3 – Link-State database Request
  - 4 – Link-State Database Update
  - 5 – Link-State Database Acknowledgment

Keep in mind that OSPFv3 message header fields differ in that there:

- Are no fields for authentication
- Is an instance ID field that has local link significance only
- The mechanisms for neighbor discovery and adjacency formation have not changed.
- The supported interface types point-to-point, point-to-multipoint, broadcast, NBMA, and virtual have not changed.
- LSA flooding and aging have not changed.
- All of OSPFv2 optional capabilities, including on-demand circuit support, NSSA areas, and the multicast extensions to OSPF are supported in OSPFv3.
- Area ID and Router ID remain 32 bit identifiers. Areas can be configured for Not-So-Stubby-Area (NSSA), Stub Area, and virtual-links.

## IPsec for OSPFv3

IPsec authentication and encrypted authentication are supported. The IPsec authentication algorithms supported are:

- Message-Digest algorithm 5 (**MD5**)
- Secure Hash Algorithm 1 (**SHA1**)
- Advanced Encryption Standard Cipher Algorithm in Cipher Block Chaining (**AESCBC**)

The IPsec encryption algorithms supported are:

- Triple Data Encryption Standard (**3DES**)
- **AESCBC** (with 128, 192, or 256 bit keys)

## Implementing OSPFv3

To implement OSPFv3 in your network:

- Map out the AS including routers and the areas to which they belong
- Create an OSPFv3 routing instance for this AS
- Configure each router with a router ID
- Configure the area that each router belongs to
- Enable OSPFv3 on each routing interface for the router specifying the OSPFv3 process, area and optional instance
- Optionally enable IPsec and configure IPsec authentication or encrypted authentication on each routing interface
- Optionally determine which router will be the designated router and backup and configure OSPF priority values on each routing interface accordingly
- Optionally configure OSPFv3 timers
- Optionally, configure the protocols and route types that will be redistributed over this AS
- Optionally configure interface cost
- Optionally modify the administrative distance for OSPF routes
- Optionally enable graceful restart
- Optionally enable the BFD protocol on all OSPF interfaces
- Optionally enable the PE-CE protocol on the router
  - With the PE-CE protocol enabled, optionally, configure a primary or secondary domain ID for this router
  - Optionally redistribute BGP discovered routes over OSPF

## OSPFv3 Configuration Overview

OSPFv3 is enabled by creating an OSPFv3 process using the **ipv6 router ospf** command in OSPFv3 router configuration mode. Once a process is created, the router's OSPF settings are configured with respect to the process ID and IP interfaces. By default, OSPFv3 is disabled on the 7100-Series device. Be aware that unspecified parameters use their default values, and any parameters specified at the interface level will override the values specified at the area level for that interface.

### Configuring Basic OSPFv3 Parameters

Basic OSPFv3 configuration consists of:

- Creating an OSPFv3 routing instance on the router
- Entering interface configuration mode for the routing interfaces for this device
- Enabling OSPFv3 on each routing interface, specifying the OSPFv3 process, area, and optional instance
- Enabling the interface



## Configuring a Routing Instance

OSPFv3 routing configuration takes place within a routing instance. Configure a routing instance using the **ipv6 router ospf** command in global configuration command mode. Executing this command places you in the OSPFv3 router configuration command mode for the specified OSPFv3 router instance.

The following example creates an OSPFv3 routing instance 23 on the router:

```
7100 Chassis(su)->configure
7100 Chassis(su-config)->ipv6 router ospf 23
7100 Chassis(su-config-ospfv3)->
```

## The IPv6 Link-Local Address

OSPFv3 uses the IPv6 link-local interface addresses as both the source and next-hop addresses. A single link-local address is supported per interface. IPv6 link-local addresses begin with FF80::/10. An IPv6 link-local address is autogenerated when you enable IPv6 on the interface, based upon the the MAC address associated with the interface. Use the **ipv6 enable** command to enable IPv6 on a non-routing IPv6 interface. Use the **ipv6 forwarding** command to enable IPv6 on a routing interface. Use the **ipv6 address link-local** option to manually configure the IPv6 link-local address. When manually configuring a link-local address, if a link-local address already exists on the interface, a warning displays asking you if you wish to change it.

Enable the interface using the **no shutdown** command.

## Configuring OSPFv3 on the Routing Interface

OSPFv3 is enabled on an IPv6 routing interface using the **ipv6 ospf** command, specifying the OSPFv3 process and the OSPF area to which the interface belongs. A basic OSPFv3 routing interface configuration does not require a configured IPv6 address, only the link-local. If you want to configure an IPv6 address on the interface, use the **ipv6 address** command to manually configure an IPv6 address on the interface or the **ipv6 address autoconfig** command to have an IPv6 address autogenerated for the interface. The following OSPFv3 basic routing interface example:

- Creates interface VLAN 1255
- Enables IPv6 forwarding on the interface
- Displays a summary of the the IPv6 interface configuration
- Enables OSPFv3 for process 23 area 0.0.0.1 on the interface

```
7100 Chassis(su)->configure
7100 Chassis(su-config)->interface vlan 1255
7100 Chassis(su-config-intf-vlan.0.1255)->ipv6 forwarding
7100 Chassis(su-config-intf-vlan.0.1255)->no shutdown
7100 Chassis(su-config-intf-vlan.0.1255)->show ipv6 interface vlan.0.1255 brief
Oper Status Legend: up = up; dn = down; tn = tentative dp = duplicate


--Status--
Interface      IPv6 Address                               Pfx Adm Op Fwd
-----
vlan.0.1255    fe80::21f:45ff:fe9a:6cb7%1255             10 dn  dn yes

7100 Chassis(su-config-intf-vlan.0.1255)->ipv6 ospf 23 area 0.0.0.1
7100 Chassis(su-config-intf-vlan.0.1255)->
```

## Basic OSPF Topology

Figure 33-1 provides an overview of a basic OSPFv3 topology. This topology displays two areas: a backbone area which must exist in any OSPF topology and a directly connected area 1. See “Configuring OSPFv3 Areas” on page 33-11 for a full discussion of OSPF area configuration. This basic configuration requires the configuration of three interfaces and two routers on the OSPFv3 router instance. Because OSPFv3 uses link-local addresses for source and next-hop, an IPv6 address is not required for a minimal OSPF topology. The link-local address is autogenerated on the interface when IPv6 forwarding is enabled. For OSPFv3, the router ID is used as the neighbor ID. For all other OSPFv3 configuration, default values are used.

**Figure 33-1 Basic OSPF Topology**



- |  |  |
|--|--|
| <ol style="list-style-type: none"> <li>1 VLAN 1, link-local address:<br/>fe80::20a:45af:fe9a:6ca5%1</li> <li>2 VLAN 2, link-local address:<br/>fe80::38c:53fd:fa8a:3bb2%2</li> </ol> | <ol style="list-style-type: none"> <li>3 VLAN 2, link-local address:<br/>fe80::45f:43ff:fa9b:3bb2%2</li> <li>4 VLAN 3, link-local address:<br/>fe80::a9b:26fa:fc4a:6fa1%3</li> </ol> |
|--|--|

## Example

The following example configures the basic OSPF topology as displayed in Figure 33-1 on page 33-7:

### Router 1 CLI Input

```
Router 1 (rw)->configure
Router 1 (rw-config)->interface vlan 1
Router 1 (rw-config-intf-vlan.0.1)->ipv6 forwarding
Router 1 (rw-config-intf-vlan.0.1)->ipv6 ospf 1 area 0.0.0.1
Router 1 (rw-config-intf-vlan.0.1)->no shutdown
Router 1 (rw-config-intf-vlan.0.1)->exit
Router 1 (rw-config)->interface vlan 2
Router 1 (rw-config-intf-vlan.0.1)->ipv6 forwarding
Router 1 (rw-config-intf-vlan.0.1)->ipv6 ospf 1 area 0.0.0.1
Router 1 (rw-config-intf-vlan.0.1)->no shutdown
Router 1 (rw-config-intf-vlan.0.2)->exit
```

```
Router 1 (rw-config) ->ipv6 router ospf 1
Router 1 (rw-101-config-ospfv3) ->router-id 1.1.1.1
Router 1 (rw-101-config-ospfv3) ->exit
Router 1 (rw-config) ->
```

## Router 2 CLI Input

```
Router 2 (rw) ->configure
Router 2 (rw-config) ->interface vlan 2
Router 2 (rw-config-intf-vlan.0.2) ->ipv6 forwarding
Router 2 (rw-config-intf-vlan.0.2) ->ipv6 ospf 1 area 0.0.0.1
Router 2 (rw-config-intf-vlan.0.2) ->exit
Router 2 (rw-config) ->interface vlan 3
Router 2 (rw-config-intf-vlan.0.3) ->ipv6 forwarding
Router 2 (rw-config-intf-vlan.0.3) ->ipv6 ospf 1 area 0.0.0.0
Router 2 (rw-config-intf-vlan.0.3) ->exit
Router 2 (rw-config) ->ipv6 router ospf 1
Router 2 (rw-config-ospfv3) ->router-id 2.2.2.2
Router 2 (rw-config-ospfv3) ->exit
Router 2 (rw-config) ->
```

## Configuring the Router ID

In OSPFv3, all neighbors are identified by their router ID, not the interface address, as is the case for OSPFv2 broadcast and NBMA links. OSPFv3 initially assigns all routers a router ID based on the highest loopback IP address of the interfaces configured for IP routing. If there is no loopback interface configured then it will be the highest VLAN IP address configured. This unique value, which is included in the hello packet transmitted in Link State Advertisements (LSA), identifies one router to another and helps establish adjacencies among OSPFv3 routers. When you specify an interface as the router ID, this value supersedes the default ID.

Use the **router-id** command in OSPFv3 configuration mode to override the default router ID.

The following example sets the router ID to 1.1.1.1, overriding the default router ID value:

```
7100-Series (rw) ->configure
7100-Series (rw-config) ->interface loopback 1
7100-Series (su-config-intf-loop.0.1) ->ip address 2001:a123::1
7100-Series (rw-config-intf-loop.0.1) ->exit
7100-Series (rw-config) ->ipv6 router ospf 1
7100-Series (rw-config-ospfv3) ->router-id 1.1.1.1
7100-Series (rw-config-ospfv3) ->exit
7100-Series (rw-config) ->
```

## Configuring the Designated Router

In the process of implementing OSPFv3, a large number of multi-access links to routers across the network may cause too many adjacencies to form. To avoid this problem, a Designated Router (DR) is elected per multi-access network to build adjacencies to all other routers on that network.

A Backup Designated Router (BDR) is also elected in case the Designated Router (DR) fails, in which case the BDR will become the DR.



**Note:** A DR is required only for multi-access networks. Point-to-Point links do not need a DR because only a single adjacency is required.

To elect a DR from a host of candidates on the network, each router multicasts an hello packet and examines the priority of hello packets received from other routers. The router with the highest priority is elected the DR, and the router with the next highest priority is elected the BDR. Any router with a priority of 0 will opt out of the DR election process. See the “[Configuring Router Priority](#)” on page 33-9 for details on configuring router priority. If DR candidates all share the same non-zero priorities, OSPF applies the router ID as a tie-breaker where the highest ID is chosen DR and the next highest ID is chosen BDR. If the priorities are not same, router ID values are not used and the highest priority determines the DR and BDR.

## Configuring Router Priority


When two routers attached to a network both attempt to become the designated router, the one with the highest router priority takes precedence. A router whose router priority is set to 0 is ineligible to become the designated router on the attached network. Router priority is specified per router interface and is advertised in hello packets sent out by the interface.

Use the **ipv6 ospf priority** command in interface configuration command mode to specify the router priority that will be included in LSAs going out this interface.

[Figure 33-2](#) on page 33-10 displays a designated router topology example. The example will configure the four displayed routers with the following priorities:

- Router 1 = 25
- Router 2 = 10
- Router 3 = 30
- Router 4 = 0

Router 4 will not take part in the election process at all. Router 3 has the highest priority and therefore will be elected DR. Router 1 has the second highest priority and will be elected BDR.

**Figure 33-2 OSPF Designated Router Topology**

### Example

The following example provides the input required to configure the designated router topology as displayed in [Figure 33-2](#) on page 33-10:

#### Router 1

```
Router 1 (rw) ->configure
Router 1 (rw-config) ->interface vlan 1
Router 1 (rw-config-intf-vlan.0.1) ->ipv6 ospf priority 25
Router 1 (rw-config-intf-vlan.0.1) ->exit
Router 1 (rw-config) ->
```

#### Router 2

```
Router 2 (rw) ->configure
Router 2 (rw-config) ->interface vlan 1
Router 2 (rw-config-intf-vlan.0.1) ->ipv6 ospf priority 10
Router 2 (rw-config-intf-vlan.0.1) ->exit
Router 2 (rw-config) ->
```

#### Router 3

```
Router 3 (rw) ->configure
Router 3 (rw-config) ->interface vlan 1
Router 3 (rw-config-intf-vlan.0.1) ->ipv6 ospf priority 30
Router 3 (rw-config-intf-vlan.0.1) ->exit
Router 3 (rw-config) ->
```

#### Router 4

```
Router 4 (rw) ->configure
Router 4 (rw-config) ->interface vlan 1
Router 4 (rw-config-intf-vlan.0.1) ->ipv6 ospf priority 0
```

```
Router 4 (rw-config-intf-vlan.0.1)->exit
Router 4 (rw-config)->
```

## Configuring the Administrative Distance for OSPF Routes

If several routes coming from different protocols are presented to the Route Table Manager (RTM), the protocol with the lowest administrative distance will be chosen for route installation. The 7100-Series platform supports connected, static, OSPF, and RIP routes.

The table below displays the default distance for these routing protocols.

Route Source	Default Distance
Connected	0
Static	1
OSPF	110
RIP	120

Use the **distance ospf** command in OSPFv3 router configuration command mode to change the administrative distance assigned to the OSPFv3 protocol. This command provides for the configuration of separate values for OSPFv3 external and intra-area routes.

## Configuring OSPFv3 Areas

OSPFv3 allows collections of contiguous networks and hosts to be grouped together. Such a group is called an area. Each area runs a separate copy of the basic link-state routing algorithm. This means that each area has its own link-state database.

The topology of an area is invisible from outside of the area, and routers internal to a given area know nothing of the detailed topology external to the area. This isolation of area detail enables the protocol to effect a marked reduction in routing traffic as compared to treating the entire Autonomous System as a single link-state domain. A router has a separate link-state database for each area it is connected to. Routers connected to multiple areas are called Area Border Routers (ABR). Two routers belonging to the same area have, for that area, identical area link-state databases.

An autonomous system can have one or more areas. A multiple area AS must define one of the areas as the backbone with an area ID of **0**. Area IDs are assigned when enabling OSPFv3 on the interface using the **ipv6 ospf** command (see [“Configuring OSPFv3 on the Routing Interface”](#) on page 33-6). All non-backbone areas in a multiple area AS must either be contiguous to the backbone or connected using a virtual-link. The backbone is responsible for distributing routing information between non-backbone areas. The backbone must be contiguous. However, it need not be physically contiguous; backbone connectivity can be established and maintained through the configuration of virtual-links.

Virtual-links can be configured between any two backbone routers that have an interface to a common non-backbone area. Such virtual-links belong to the backbone. The protocol treats two routers joined by a virtual-link as if they were connected by an unnumbered point-to-point backbone network.

Inter-area route calculation for OSPFv3 is similar to OSPFv2. See RFC 2328 for inter-area route calculation details. See RFC 2740 *OSPF for IPv6* for details on the differences between calculating inter-area routes between OSPFv2 and OSPFv3.

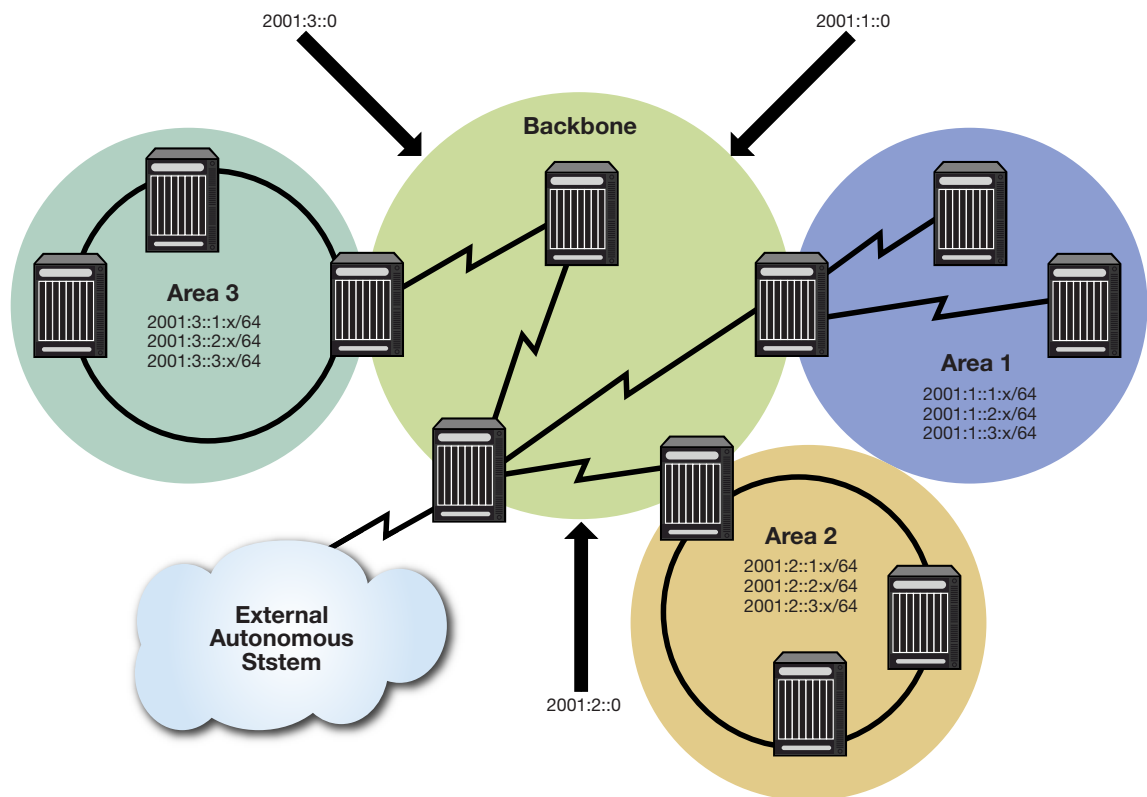
An Area ID can be any value from 0 - 4294967295, but is converted into the 32-bit dotted-quad format (area 50 would be displayed as 0.0.0.50; area 3546 would be displayed as 0.0.13.218)

## Configuring Area Range

An area range is a form of address summarization that defines a range of addresses to be used by the backbone ABRs when they communicate routes to other areas. Area range is a critical tool that pares the route tables and update traffic, as well as reduces network recalculation by the Dijkstra algorithm. Area range configuration summarizes by aggregating an areas' internal addresses to advertise a single network. Backbone routers see only one update, representing an entire range of subnets. Area ranges can be configured for purposes of advertisement as well as summarization of addresses that should not be advertised.

Use the **area range** command in OSPFv3 configuration command mode to configure an area address summarization.

**Figure 33-3 OSPF Summarization Topology**



## Example

The following example provides the input required to configure summarization of the three area topology as displayed in [Figure 33-3](#) on page 33-12:

### Area 1

```
ABR1 (rw) ->configure
ABR1 (rw-config) ->router ospf 1
ABR1 (rw-config-ospf-1) ->area 1 range 2001:1::0/64
ABR1 (rw-config-ospf-1) ->exit
ABR1 (rw-config) ->
```

## Area 2

```
ABR2 (rw) ->configure
ABR2 (rw-config) ->router ospf 1
ABR2 (rw-config-ospf-1) ->area 2 range 2001:2::0/64
ABR2 (rw-config-ospf-1) ->area 2 range 2001:2:4::0/64 not-advertised
ABR2 (rw-config-ospf-1) ->exit
ABR2 (rw-config) ->
```

## Area 3

```
ABR3 (rw) ->configure
ABR3 (rw-config) ->router ospf 1
ABR3 (rw-config-ospf-1) ->area 3 range 2001:3::0/64
ABR3 (rw-config-ospf-1) ->exit
ABR3 (rw-config) ->
```

## Configuring a Stub Area

A stub area is a non-transit area. In other words, an area that does not originate or propagate external routes. AS-external-LSAs are not flooded into the stub area; routing to AS external networks is based on a single per-area default route. This reduces the link-state-database size and memory requirements for routers within stub areas.

Handy for reducing routing table size, a stub area is a “dead-end” in which there is no other way to enter or exit except through an Area Border Router (ABR). No ASE (Autonomous System External) or NSSA routes are permitted in a stub area. Each router in a stub area must specify that they are members of the stub area. When specifying that the ABR is a member of the stub area, the ABR will inject a default route into the area.

Routing to external designations from stub areas is based on a default route injected by a stub area’s ABR. A default route is automatically created by the stub area’s ABR. This default route is injected into the stub area to enable other stub routers within the stub area to reach any external routes that are no longer inserted into the stub area.

A stub area can be configured such that the ABR is prevented from sending type 3 summary LSAs into the stub area using the **no-summary** option. In this case, all destinations outside of the stub area are represented by means of a default route.

There are a couple of restrictions on the use of stub areas. Virtual-links cannot be configured through stub areas, and AS boundary routers cannot be placed internal to stub areas.


Use the **area stub** command in OSPF router configuration command mode to configure an area as a stub.

### Stub Area Default Route Cost

A cost value can be set for the default route that is sent into a stub area by an ABR. Configuration of the stub area default route cost is restricted to the ABR attached to this stub area.

Use the **area default-cost** command in OSPFv3 router configuration command mode on the ABR attached to this stub area to configure the stub area default route cost.



**Figure 33-4 OSPF Stub Area Topology**

### Example

Every router in Areas 1 and 2 are configured for a stub area (Routers 1, 2, and 3 for Area 1 and Routers 5, 6, 7, and 8 for Area 2). Additionally, ABR routers 3, 5, and 6 are also configured with a default-cost to be assigned to the stub area. Router 5 has a lower metric cost when compared to Router 6, so Router 5 will be the preferred router for packets to access the area, with Router 6 employed as a backup in case Router 5 fails. The following example provides the input required to configure the stub topology as displayed in [Figure 33-4](#) on page 33-14:

#### Router 1

```
Router1(rw-config)->router ospf 1
Router1(rw-config-ospf-1)->area 1 stub
```

#### Router 2

```
Router2(rw-config)->router ospf 1
Router2(rw-config-ospf-1)->area 1 stub
```

#### Router 3

```
Router3(rw-config)->router ospf 1
Router3(rw-config-ospf-1)->area 1 stub
Router3(rw-config-ospf-1)->area 1 default-cost 15
```

#### Router 5

```
Router5(rw-config)->router ospf 1
Router5(rw-config-ospf-1)->area 2 stub
Router3(rw-config-ospf-1)->area 2 default-cost 15
```

#### Router 6

```
Router6(rw-config)->router ospf 1
Router6(rw-config-ospf-1)->area 2 stub
Router6(rw-config-ospf-1)->area 2 default-cost 20
```

#### Router 7

```
Router7(rw-config)->router ospf 1
Router7(rw-config-ospf-1)->area 2 stub
```

#### Router 8

```
Router8(rw-config)->router ospf 1
Router8(rw-config-ospf-1)->area 2 stub
```

## Configuring a Not So Stubby Area (NSSA)

A Not So Stubby Area (NSSA) is a hybrid area using an Autonomous System Border Router (ASBR) to connect two disparate organizations. External routes are advertised as Type 7 LSAs and are converted to Type 5 LSAs before flooding to the backbone by the NSSA's ABR. Also, summary routes are allowed into the NSSA while external routes from other networks are still filtered from insertion into the NSSA.


External routes that are not imported into an NSSA can be represented by a default route. If the router is an ABR and has the highest router ID of all ABRs in the area, and no other ABR in the area is configured to translate always, it will translate Type 7 LSAs into Type 5 LSAs. Configuring the identity of the translator can be used to bias the routing to aggregated destinations. When translator role is set to Always, Type-7 LSAs are always translated regardless of the translator state of other NSSA border routers.

When a translating ABR loses a translator election, it will stop translating, and after a number of seconds (set by the **transstabilityint** option), it will flush any Type 5 LSAs resulting from aggregation. Any Type 5 LSAs resulting from direct translation of Type 7 LSAs will be allowed to age out. An ABR will always originate a default route into any attached NSSAs.

If the **no-summary** option is specified, the ABR does not send type 3 summary LSAs into the NSSA area, therefore all destinations outside of the NSSA area are represented by means of a default route.

Use the **area nssa** command to configure an area as a Not-So-Stubby-Area.

**Figure 33-5 OSPF NSSA Topology**



### Example

Routers 2 and 6 are configured as the ABRs between Area 1 and 0, and Router 4 as the ASBR. Router 2 is configured to set Area 1 as an NSSA, and Type 7 routes from the connected domain will be translated to Type 5 routes into the backbone.

ABR Router 2 will only translate Type 7 LSAs; static routes redistributed by router 4. Also, Router 2 will always translate, since it is configured to do so; Router 6 will not, since only one ABR will perform the translation for a given area.

Router 4 will be configured to redistribute static routes.

The following example provides the input required to configure the NSSA topology as displayed in [Figure 33-5](#) on page 33-15:

#### Router 6 (ABR)

```
Router 6 (rw) ->configure
```

```
Router 6 (rw-config)->interface vlan 1
Router 6 (rw-config-intf-vlan.0.1)->ipv6 address 2001:0:1:::1:1/64
Router 6 (rw-config-intf-vlan.0.1)->ipv6 forwarding
Router 6 (rw-config-intf-vlan.0.1)->ipv6 ospf 1 area 0.0.0.0
Router 6 (rw-config-intf-vlan.0.1)->no shutdown
Router 6 (rw-config-intf-vlan.0.1)->exit
Router 6 (rw-config)->interface vlan 2
Router 6 (rw-config-intf-vlan.0.2)->ipv6 address 2001:1:1:::1:1/64
Router 6 (rw-config-intf-vlan.0.2)->ipv6 forwarding
Router 6 (rw-config-intf-vlan.0.2)->ipv6 ospf 1 area 0.0.0.1
Router 6 (rw-config-intf-vlan.0.2)->no shutdown
Router 6 (rw-config-intf-vlan.0.2)->exit
Router 6 (rw-config)->ipv6 router ospf 1
Router 6 (rw-config-ospfv3)->router-id 6.6.6.6
Router 6 (rw-config-ospfv3)->area 1 nssa
Router 6 (rw-config-ospfv3)->exit
```

### Router 2(ABR)

```
Router 2 (rw)->configure
Router 2 (rw-config)->interface vlan 1
Router 2 (rw-config-intf-vlan.0.1)->ipv6 address 2001:0:1:::2:1/64
Router 2 (rw-config-intf-vlan.0.1)->ipv6 forwarding
Router 2 (rw-config-intf-vlan.0.1)->ipv6 ospf 1 area 0.0.0.0
Router 2 (rw-config-intf-vlan.0.1)->no shutdown
Router 2 (rw-config-intf-vlan.0.1)->exit
Router 2 (rw-config)->interface vlan 2
Router 2 (rw-config-intf-vlan.0.2)->ipv6 address 2001:1:1:::2:1/64
Router 2 (rw-config-intf-vlan.0.2)->ipv6 forwarding
Router 2 (rw-config-intf-vlan.0.2)->ipv6 ospf 1 area 0.0.0.1
Router 2 (rw-config-intf-vlan.0.2)->no shutdown
Router 2 (rw-config-intf-vlan.0.2)->exit
Router 2 (rw-config)->ipv6 router ospf 1
Router 2 (rw-config-ospfv3)->router-id 2.2.2.2
Router 2 (rw-config-ospfv3)->area 1 nssa
Router 2 (rw-config-ospfv3)->area 1 nssa transrole always
Router 2 (rw-config-ospfv3)->area 1 nssa-range 2002:1:1:::1/64
Router 2 (rw-config-ospfv3)->exit
```

### Router 4 (ASBR)

```
Router 4 (rw)->configure
Router 4 (rw-config)->interface vlan 2
Router 4 (rw-config-intf-vlan.0.2)->ipv6 address 2001:1:1:::2:2/64
Router 4 (rw-config-intf-vlan.0.2)->ipv6 forwarding
Router 4 (rw-config-intf-vlan.0.2)->ipv6 ospf 1 area 0.0.0.1
Router 4 (rw-config-intf-vlan.0.1)->no shutdown
Router 4 (rw-config-intf-vlan.0.1)->exit
```

```
Router 4 (rw-config)->interface vlan 3
Router 4 (rw-config-intf-vlan.0.3)->ipv6 address 2001:3:1::1:1/64
Router 4 (rw-config-intf-vlan.0.3)->ipv6 forwarding
Router 4 (rw-config-intf-vlan.0.3)->ipv6 ospf 1 area 0.0.0.1
Router 4 (rw-config-intf-vlan.0.3)->no shutdown
Router 4 (rw-config-intf-vlan.0.3)->exit
Router 4 (rw-config)->ipv6 router ospf 1
Router 4 (rw-config-ospfv3)->router-id 4.4.4.4
Router 4 (rw-config-ospfv3)->redistribute static metric-type 1
Router 4 (rw-config-ospfv3)->exit
```


## Configuring Area Virtual-Links

The backbone area 0 cannot be disconnected from any other areas in the AS. Disconnected areas will become unreachable. To establish and maintain backbone connectivity, virtual-links can be configured through non-backbone areas for the purpose of connecting a disconnected area with the backbone through a backbone connected area. The two endpoints of a virtual-link are ABRs, both of which belong to the backbone connected area (also referred to as the transit area); one of which belongs to the area disconnected from the backbone. Virtual-links cannot be configured through stub areas (see “[Configuring a Stub Area](#)” on page 33-13 for stub area configuration information).

The virtual-link is treated as if it were an unnumbered point-to-point network belonging to the backbone and joining the two ABRs. The cost of a virtual-link is not configured. It is auto configured with the cost of the intra-area path between the two ABRs that make up the virtual-link.

Use the **area virtual-link** command in OSPF router configuration command mode, providing the transit area ID and the ABRs IP address, to configure an area virtual-link.

[Figure 33-6](#) on page 33-18 displays a typical virtual-link topology. Area 3 does not share an ABR with the backbone area, and is therefore disconnected from the backbone. Area 3 shares an ABR (router 2) with area 1. Area 1 has a second ABR (router 1) that it shares with the backbone. Area 1 is the transit area because it contains an ABR that it shares with the disconnected area and a second ABR that it shares with the backbone. By configuring an area virtual-link between router 2 and router 1, Area 3 will gain connectivity with the backbone and be able to learn routes for this AS.

**Figure 33-6 Virtual-Link Topology**

## Example

The following example presents the configuration required to configure the virtual-link displayed in [Figure 33-6](#) on page 33-18:

### Router 1

```
Router 1 (rw-config)->router ospf 1
Router 1 (rw-config-ospf-1)->area 0.0.0.1 virtual-link 2.2.2.2
Router 1 (rw-config-ospf-1)->exit
Router 1 (rw-config)->
```

### Router 2

```
Router 2 (rw-config)->router ospf 2
Router 2 (rw-config-ospf-2)->area 0.0.0.1 virtual-link 1.1.1.1
Router 2 (rw-config-ospf-2)->exit
Router 2 (rw-config)->
```

### Configuring Area Virtual-Link Timers

The following timers can be configured for an area virtual-link:

- Dead-interval using the **area virtual-link dead-interval** command
- Hello-interval using the **area virtual-link hello-interval** command
- Retransmit-interval using the **area virtual-link retransmit-interval** command
- Transmit-delay using the **area virtual-link transmit-delay** command

See [“Configuring OSPFv3 Timers”](#) on page 33-23 for an OSPF timers discussion.



**Note:** RFC 2328 specifies that the retransmit-interval should be greater than the expected round-trip delay between the two routers. This may be hard to estimate for a virtual-link; it is better to err on the side of making it too large.

## Configuring IPsec Authentication for OSPFv3

Internet Protocol Security (IPsec) is an internet layer, end-to-end security scheme that provides for the securing of IP communications by authentication and encrypted authentication of each communication session IP packet. IPsec for OSPFv3 is configured on the interface. The IPsec for OSPFv3 implementation supports both authentication only and encrypted authentication. IPsec must first be enabled on the interface. Supported IPsec authentication algorithms are:

- **MD5** – Message-Digest algorithm 5
- **SHA1** – Secure Hash Algorithm 1
- **AESCBC** – Advanced Encryption Standard (AES) Cipher Algorithm in Cipher Block Chaining (CBC)

Supported IPsec encryption ciphers are:

- **3DES** – Triple Data Encryption Standard cipher algorithm
- **AESCBC** – AES (Cipher Block Chaining) cipher algorithm

Each IPsec configuration must have a Security Parameters Index (SPI) with a value between **256 - 4294967295** assigned to it and a security key. The key can be specified as a hex key.

IPsec must be enabled in global VRF router configuration mode using the **crypto ipsec enable** command before using IPsec for OSPFv3 authentication or encrypted authentication.

Configure IPsec for OSPFv3 on an interface for authentication only by specifying the SPI and authentication algorithm using the **ipv6 ospf authentication** command in interface configuration mode.

This example shows how to configure VLAN 1 for IPsec SPI entry 256 for MD5 authentication with a hex key of **1234567890abcdef**:

```
7100 Chassis(rw-config)->crypto ipsec enable
7100 Chassis(rw-config)->interface vlan 1
7100 Chassis(rw-config-intf-vlan.0.1)->ipv6 ospf authentication spi 256 md5
1234567890abcdef hex
```

Configure IPsec for OSPFv3 on an interface for encrypted authentication by specifying the SPI, authentication algorithm and encryption cipher using the **ipv6 ospf encryption** command in interface configuration mode.

This example shows how to configure VLAN 1 for IPsec SPI entry 256 for the 128-bit aescbc encryption with a key of **1234567890abcdef**, and for MD5 authentication with a hex key of **1234567890abcdef**:

```
7100 Chassis(rw-config)->crypto ipsec enable
7100 Chassis(rw-config)->interface vlan 1
7100 Chassis(rw-config-intf-vlan.0.1)->ipv6 ospf encryption ipsec spi 256 esp
aescbc 128 1234567890abcdef hex auth md5 1234567890abcdef hex
```

## Configuring Route Redistribution

Redistribution permits the importation of other routing protocols into OSPF such as RIP, as well as static and directly connected routes. Alternately, you can specify a route-map for redistribution into OSPF. Be aware that if the referenced route map has not yet been configured, then an empty

route map is created with the specified name. See “[Configuring a Not So Stubby Area \(NSSA\)](#)” on page 33-15 for an example of redistribution of static routes by an ASBR in an NSSA context.

Use the **redistribute** command in OSPF router configuration command mode to permit the redistributions of OSPF, RIP, static, or connected routes by this router.

## Filtering Routes from the OSPF Route Table

Routes can be filtered from the OSPF route table by creating an OSPF filter route route-map and assigning it to the distribute-list for this OSPFv3 router.

Use the **route-map filter** command as described in the “Route-Map Manager” section of the *Extreme Networks 7100-Series CLI Reference* to create an OSPF filter route route-map.

Use the **distribute-list route-map in** command to assign the filter route-map to the OSPF distribute-list.

## Configuring Passive Interfaces

Passive interfaces explicitly allow the network to be advertised, but the router prevents the forming of neighbor relationships on that interface. Passive interfaces are included in the OSPF route table. Passive interfaces do not send or receive hello packets. OSPF adjacencies can not be formed on a passive interface.

An option exists to default all interfaces to passive mode.

Use the **passive-interface** command in router configuration command mode to configure an interface as passive.

## Graceful Restart

OSPF graceful restart, sometimes referred to as non-stop forwarding, provides for an OSPF router to remain on the forwarding path during a restart of its OSPF software. Graceful restart has three elements to its configuration: enabling, helper router, and restart interval.

Enabling graceful restart instructs the firmware to perform a graceful restart, rather than a standard OSPF restart. Restart is only initiated by a fail-over. Grace LSAs are sent when OSPF is restarted on another module. Whether the failover is intentional or not, the failed router protocol is restarted on another module, and upon startup, OSPF sends grace LSAs out to its neighbors using existing link aggregation groups. Use the **graceful-restart enable** command to enable the graceful restart ability on this router.

The helper relationship with the restarting router is on a per network segment basis. The helper monitors the network for topology changes. If no changes occur, the helper router continues to advertise its LSAs as though no restart was occurring. If the restarting router was the designated router, the helper continues to treat it as such. If a topology change does occur, graceful restart is terminated on the restarting router and a standard restart occurs. Helper mode is enabled by default. Helper mode can be disabled on a restarting router neighbor using the **ipv6 ospf helper-disable** command in interface command mode. If the restarting router receives an LSA indicating a disabled helper, the graceful restart terminates and a standard restart occurs.

A restart interval provides for a maximum time in seconds after which the graceful restart will terminate should it not complete or terminate for other reasons within the interval. Use the **graceful-restart restart-interval** command to change the restart interval setting.

View the router OSPF section of the **show running-config** display to verify any non-default graceful restart settings.

## Graceful Restart and High Availability

7100-Series modules support single router high availability failover using the following components:

- OSPF graceful restart
- Non-stop router frame forwarding on each module
- Single router configuration
- Router protocol process failover to another module
- Link Aggregate Group (LAG) connectivity to neighboring routers

In a stable network, the route and rule information is fairly constant. If the router protocol process was to suddenly fail, forwarding information current at the time of the failure in all probability is usable for the short time after the failure until recovery occurs. During this recovery period, existing connections (that were not directly using the failed module) remain in effect. New connections continue to be installed using the last known “good” forwarding information. The router protocol process that failed is dynamically restarted. You do not configure where the router process is running. The router forwarding process remains active on every module. The protocol process exchanges protocol and maintains state that it distributes to the other modules and does not have to run on any specific module. One exception to this rule is that the module must have 256M of memory to be router protocol process eligible.

Upon failure of a module running the router protocol process, the protocol process is started on a recovery module. One of the first messages it sends to its OSPF neighbors is a grace LSA. High availability failover will successfully occur if the following is true:

- The router is enabled for graceful restart
- The neighbors are enabled to participate as graceful restart helpers
- The OSPF dead interval is configured for a sufficient period such that the grace LSA is received by its neighbors before the configured OSPF dead interval expires
- And each neighbor is a member of a LAG common to the failed router, allowing the neighbor to remain up




**Figure 33-7 Physical and Logical Single Router HA Failover Configuration**

Figure 33-7 depicts the physical and logical configurations of the single router high availability failover mechanism. The neighbor to router lines display direct neighbor connections to the router enabled for OSPF graceful restart and members of LAGs common to the failing router. See [Chapter 21, Link Aggregation Control Protocol \(LACP\) Configuration](#) for LAG configuration details. The server 100.1.1.3 and 100.1.1.5 to router lines display VLAN connections common to both the failing and recovery routers. Helper mode on each neighbor is enabled by default. Enable graceful restart on the router using the **graceful-restart enable** command.

```
7100 Chassis(rw-config)->router ospf 1
7100 Chassis(rw-config-ospf-1)->graceful-restart enable
7100 Chassis(rw-config-ospf-1)->
```

## Configuring Interface Cost

Each interface has an outbound cost associated with it. The lower the cost, the more likely the interface will be used to forward data traffic. Should several equal-cost routes to a destination exist, traffic is distributed equally among them.

The formula for calculating the OSPF interface cost metric is the reference bandwidth divided by the interface bandwidth. By default the reference bandwidth is set to 100 Mbps. For 10 Mbps links, the resulting cost is 10. For 100, 1000, or 10000 Mbps links, the resulting cost is 1. The reference bandwidth can be modified using the **auto-cost reference-bandwidth** command in OSPF configuration mode. The ability to re-center the reference bandwidth to a higher value, allows for OSPF interface costs to default to a value greater than 1 for 100, 1000, or 10000 Mbps links and greater than 10 for 10 Mbps links.

It is recommended that the auto cost reference bandwidth be the same value for all OSPF routers in the domain.

Use the **ipv6 ospf cost** command in interface configuration command mode to statically specify the outbound cost of this interface. A statically configured OSPF interface cost overrides all other interface cost methods.

For logical interfaces containing multiple physical interfaces, such as a LAG, the aggregate interface speed is not readily available. A tracked object configured with the ports belonging to the logical interface can return the physical interface speed of each physical port specified in the tracked object. OSPF will sum the returned interface speeds and use that aggregate value when calculating OSPF interface cost. Because the tracked object will report when a physical interface is up or down, OSPF will dynamically adjust the aggregate speed when an interface becomes active or goes down and adjust the OSPF interface cost accordingly. This method should be used in LAG and ECMP logical interface contexts.



**Note:** The speed used in the cost calculation is sum of all ports capabilities in the tracked object. Setting the speed manually will not change the tracked interface speed. A 1GB capable port has a 1 GB speed regardless of the manual speed setting. The same holds true for ports that auto-negotiate to a lower speed. The expectation is that both sides of the link are using the same ports and SFP connectors and should result in the same speed.

Use the **ipv6 ospf cost track** command in interface configuration mode to calculate the OSPF interface cost based upon summing physical interface speeds that belong to a logical interface.

When adding an additional physical port to a logical interface that uses the interface summation method to determine OSPF interface cost, you must also add the physical port to the associated tracked object.

See “[Tracked Object Manager Configuration](#)” on page 12-1 for tracked object configuration details.

## Configuring Bidirectional Forwarding Detection (BFD) on Interfaces

BFD is used to detect a communications failure with an OSPF forwarding plane next-hop. BFD detects failures in under one second. BFD augments the OSPF Hello mechanism. The OSPF Hello interval defaults to 10 seconds. With high speed data rates, a failure requiring multiple seconds to detect can result in significant data loss. The OSPF implementation of the BFD protocol uses the following non-configurable parameters:

**Transmit Interval** – The period of time between the transmission of BFD control packets, set for 100ms.

**Receive Interval** – The period of time between received BFD control packets, set for 100ms.

**Detection Multiplier** – The Number of consecutive control packets that can be missed before the BFD session transitions to down, set to 3.

Use the **bfd all-intfs-on** command in OSPF router configuration mode to enable BFD on all OSPF interfaces.

## Configuring OSPFv3 Timers

There are five OSPF timers:

- Hello-Interval
- Dead-Interval
- Retransmit-Interval
- Transmit-Delay
- SPF-Delay

To ensure efficient adjacency between OSPF neighbors, the 7100-Series device provides hello-interval and dead-interval commands. The hello interval is the period between transmissions of hello packet advertisements. The dead interval is the period that can elapse without receiving a router's hello packets before its neighbors will declare it down.

Use the **ipv6 ospf hello-interval** command in interface configuration command mode to configure the period between transmissions of hello packet advertisements.

Use the **ipv6 ospf dead-interval** in interface configuration command mode to configure the period between receiving hello packets before the associated neighbor is declared down.

In order to ensure that flooding is reliable, LSAs are retransmitted until they are acknowledged. The period between retransmissions is the retransmit-interval. If this interval is set too low for an interface, needless retransmissions will take place. If the value is set too high, the speed of the flooding, during the period of lost packets, may be affected.

Use the **ipv6 ospf retransmit-interval** command in interface configuration command mode to configure the retransmit-interval.

The transmit-delay is an estimation of the number of seconds it takes to transmit a link state update packet over this interface. This value should take into account transmission and propagation delays.

Use the **ipv6 ospf transmit-delay** command in interface configuration command mode to configure the transmit-delay.

The SPF-delay is the amount of time that transpires between the receipt of an OSPF update and the SPF calculation.

Use the **timers spf** command in OSPFv3 router configuration command mode to specify the amount of time between receiving an OSPF update and an SPF calculation occurring.

The OSPF timers can also be configured for an area virtual-link. See "[Configuring Area Virtual-Links](#)" on page 33-17.

## Configuring the PE-CE Protocol

The PE-CE protocol allows a service provider offering Virtual Private Network (VPN) services to their customers to peer Customer Edge (CE) routers with their Provider Edge (PE) routers. RFC 6565 defines how the PE-CE protocol is implemented using the OSPF routing protocol.

When the PE router becomes a routing peer of the CE router, the PE router learns the routes that lead to the CE's site and can redistribute those routes to other PE routers that attach to the same VPN.

Use the **enable-pe-ce** command in OSPFv3 configuration mode to enable the PE-CE protocol on the router.

## The OSPF VRF Domain ID

If the OSPF instances of an OSPF domain are given one or more domain IDs, OSPF can determine whether an OSPF-originated VPN-IPv6 route belongs to the same domain as a given OSPF instance and whether the route should be redistributed to that OSPF instance as an inter-area route or as an OSPF AS-external route.

If two OSPF instances with a domain ID configured are in the same OSPF domain, the PE-CE protocol requires that the primary domain ID of the second instance must be one of the domain IDs of the first instance (either primary or secondary). If two OSPF instances with a domain ID configured are not in the same OSPF domain, the primary domain ID of each instance must not be configured as a domain ID of the other OSPF instance.

The PE-CE protocol must be enabled using the **enable-pe-ce** command to set the OSPF VRF domain ID.

Use the domain-id command in OSPFv3 configuration mode to set the domain ID for this PE router.

## Redistribute BGP into OSPF

OSPF supports the redistribution of BGP discovered routes into OSPF. A **global** option is available for the redistribution of BGP prefixes from the global router. When using the **global** option, VPN-IPv6 address prefixes are appropriately translated.

Use the **redistribute bgp** command in OSPFv3 configuration mode to redistribute BGP learned routes to other PE routers in this VPN.

## OSPF Sham Link

If a VRF contains both an OSPF-distributed route and a VPN-IPv6 route for the same IPv6 prefix, then the backdoor OSPF-distributed route is preferred over the VPN backbone route, unless the next hop interface for an installed (OSPF distributed) route is the sham link, in which case, the VPN backbone VPN-IPv6 route is used.

If it is desired to have OSPF prefer the routes through the VPN backbone over the routes through the OSPF backdoor link, then the routes through the backbone must appear to be intra-area routes. The sham link provides this appearance of an intra-area link connecting the two PE routers.

Use the **area sham-link** command in OSPFv3 configuration mode, specifying both the source and destination link addresses to configure an OSPF sham link between two VPN PE routers.

# OSPFv3 Configuration Details

This section provides details for the configuration of OSPFv3 on 7100-Series platforms.

## Default Settings

[Table 33-2](#) lists OSPF parameters and their default values.

**Table 33-2 Default OSPF Parameters**

Parameter	Description	Default Value
router ID	Provides for the identification of one router to another and helps establish adjacencies among OSPF routers.	highest loopback IP address or highest configured VLAN IP address if no loopback configuration exists
interface cost	An outbound interface value used in determining which routing interface should forward when more than one routing interface is available.	10
interface priority	A value placed on the interface that helps in determining which router will be elected designated router.	1
interface network type	Specifies the type of network an interface is connecting to.	broadcast

**Table 33-2 Default OSPF Parameters (continued)**

Parameter	Description	Default Value
LSA Thresholds	Specifies:  The number of LSA updates that force a full routing calculation  The number of LSA updates that interrupt and restart a full routing calculation  The number of LSA inter-area/external updates that force a full routing calculation  The number of intra updates that force a full routing calculation	<ul style="list-style-type: none"> <li>• 4294967295 Update starts</li> <li>• 4294967295 Update restarts</li> <li>• 50 Inter-area/external updates</li> <li>• 0 Intra updates</li> </ul>
LSA Pause Frequency	Specifies the number of cpu units SPF calculation runs before pausing.	10000
SPF delay timer	Specifies the amount of time between receiving an OSPF update and the start of an SPF calculation.	5 seconds
retransmit interval	A timer that determines the retransmission of LSAs in order to ensure reliable flooding.	5 seconds
transmit delay	Specifies the number of seconds it takes to transmit a link state update packet over this interface.	1 second
hello interval	The period between transmissions of hello packet advertisements.	<ul style="list-style-type: none"> <li>• 10 seconds for broadcast and point-to-point networks</li> <li>• 30 seconds for non-broadcast and point-to-multipoint networks</li> </ul>
dead interval	The period that can elapse without receiving a router's hello packets before its neighbors will declare it down.	40 seconds
distance	Specifies the administrative distance for OSPF routes. The available protocol with the lowest administrative distance is chosen for this route.	connected: 0 static: 1 OSPF = 110 RIP = 120
graceful restart	Provides for an OSPF router to remain on the forwarding path during a restart of its OSPF software.	disabled
graceful restart restart interval	Specifies the maximum time in seconds after which the graceful restart will terminate should it not complete or terminate for other reasons within the interval.	120 seconds
PE-CE Protocol	Enables the Customer Edge (CE) router as a Provider Edge (PE) router peers.	disabled

Procedure 33-1 describes how to configure basic OSPF parameters.

### Procedure 33-1 Configuring Basic OSPFv3 Parameters

Step	Task	Command(s)
1.	Create an OSPF routing instance in router configuration mode.	<b>ipv6 router ospf</b> <i>process-id</i>
2.	In interface configuration mode, enable OSPFv3 on each routing interface, specifying the OSPFv3 process, area, and optional instance	<b>ipv6 ospf</b> <i>process area area [instance instance-id]</i>
3.	Optionally, in interface configuration mode, configure an IPv6 address for all routing interfaces in the AS.	<b>ipv6 address</b> { <i>link-local-address link-local   ipv6-address/length   ipv6-prefix/length eui-64   autoconfig   general-prefix sub-bits/length</i> }

Table 33-3 describes how to configure basic OSPFv3 parameters.

### Table 33-3 Configuring OSPFv3 General Optional Parameters

Task	Command(s)
Optionally, change the OSPFv3 router ID for this device.	<b>router-id</b> <i>router-id</i>
Optionally, enable the OSPF PE-CE protocol.	<b>enable-pe-ce</b>
Optionally, configure the OSPF VRF domain ID.	<b>domain-id [secondary] type type value value</b>
Optionally, change the SPF LSA thresholds for this router.	<b>spf lsa-thresholds</b> <i>num-start num-restart num-intra-full num-ia-ext-full</i>
Optionally, change the SPF pause frequency to specify the number of units SPF calculation runs before pausing.	<b>spf pause-frequency</b> <i>units</i>
Optionally, change the delay, in milliseconds, between the receipt of an update and the beginning of the SPF execution.	<b>timers spf</b> <i>spf-delay</i>
Optionally, enable BFD on all OSPF interfaces.	<b>bfd all-intfs-on</b>
Optionally, change the administrative distance for OSPFv3 routes.	<b>distance [ospf {external   intra-area}] weight</b>
Optionally, define the range of addresses used by this Area Border Router (ABR) when communicating routes to other areas.	<b>area area-id range</b> <i>ipv6-address [not-advertise]</i>
Optionally, configure an area as a stub area.	<b>area area-id stub [no-summary]</b>
Optionally, set the cost for the default route that is sent into a stub area by an ABR.	<b>area area-id default-cost</b> <i>cost</i>
Optionally, configure an area as a not so stubby area.	<b>area {area-id   A.B.C.D} nssa [no-summary] [transstabilityint seconds] [transrole always]</b>
Optionally, configure an Autonomous System Border Router (ASBR) to summarize Type 7 to Type 5 routes matching the specified address and mask.	<b>area {area-id   A.B.C.D} nssa-range</b> <i>ipv6-address [not-advertise]</i>
Optionally, configure an OSPF sham link between two PE routers.	<b>area area-id sham-link</b> <i>source-ip-address destination-ip-address</i>

**Table 33-3 Configuring OSPFv3 General Optional Parameters (continued)**

<b>Task</b>	<b>Command(s)</b>
Optionally, configure an OSPF sham link authentication key password.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>authentication-key</b> <i>password</i>
Optionally, modify the OSPF sham link dead interval.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>dead-interval</b> <i>seconds</i>
Optionally, modify the OSPF sham link hello interval.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>hello-interval</b> <i>seconds</i>
Optionally, configure the OSPF sham link keychain.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>keychain</b> <i>name</i>
Optionally, specify an OSPF sham link message digest key and MD5 authentication key.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>message-digest-key</b> <i>digest-key</i> <b>md5</b> <i>auth-key</i>
Optionally, modify the OSPF sham link retransmit interval.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>retransmit-interval</b> <i>seconds</i>
Optionally, modify the OSPF sham link transmit delay period.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>transmit-delay</b> <i>seconds</i>
Optionally, modify the OSPF sham link cost.	<b>area</b> <i>area-id</i> <b>sham-link</b> <i>source-ip-address</i> <i>destination-ip-address</i> <b>cost</b> <i>cost</i>
Optionally, configure an OSPF virtual-link, which represents a logical connection between the backbone and a non-backbone connected OSPF area.	<b>area</b> <i>area-id</i> <b>virtual-link</b> <i>router-id</i> <b>area</b> <i>area-id</i> <b>virtual-link</b> <i>router-id</i> <b>dead-interval</b> <i>seconds</i> <b>area</b> <i>area-id</i> <b>virtual-link</b> <i>router-id</i> <b>hello-interval</b> <i>seconds</i> <b>area</b> <i>area-id</i> <b>virtual-link</b> <i>router-id</i> <b>retransmit-interval</b> <i>seconds</i> <b>area</b> <i>area-id</i> <b>virtual-link</b> <i>router-id</i> <b>transmit-delay</b> <i>seconds</i>
Optionally, change the bandwidth reference setting used for calculating interface cost for this OSPF instance.	<b>auto-cost</b> <b>reference-bandwidth</b> <i>bandwidth-multiplier</i>
Optionally, enable passive OSPF on the specified interface.	<b>passive-interface</b> { <i>vlan</i> <i>vlan-id</i>   <i>interface-name</i>   <b>default</b> }
Optionally, allow routing information discovered through non-OSPF protocols to be distributed in OSPF update messages.	<b>redistribute</b> { <b>connected</b>   <b>rip</b>   <b>static</b>   <b>blackhole</b> } [ <b>route-map</b> <i>name</i> ] [ <b>metric</b> <i>metric-value</i> ] [ <b>metric-type</b> <i>type-value</i> ] [ <b>tag</b> <i>tag</i> ]
Optionally, assign an OSPF route filter route-map to the OSPF distribute-list.	<b>distribute-list</b> <b>route-map</b> <i>name</i> <b>in</b>
Optionally, enable adjacency logging on this OSPFv3 router.	<b>log-adjacency</b>
Optionally, enable the graceful restart feature on this router.	<b>graceful-restart</b> <b>enable</b>
Optionally, change the graceful restart restart interval for this router.	<b>graceful-restart</b> <b>restart-interval</b> <i>interval</i>

**Table 33-3 Configuring OSPFv3 General Optional Parameters (continued)**

Task	Command(s)
Optionally, in system command mode, reset the specified OSPFv3 process ID or the OSPFv3 process.	<b>clear ipv6 ospf process</b> [ <i>process-id</i> ]
Optionally, in global configuration command mode, enable OSPFv3 protocol debugging output for the specified subsystem.	<b>debug ipv6 ospf</b> { <i>subsystem</i> }

Table 33-4 describes how to configure optional OSPF interface parameters. All commands in this procedure are entered in interface configuration command mode.

**Table 33-4 Configuring OSPF Optional Interface Parameters**

Task	Command(s)
Optionally, configure the OSPFv3 router neighbors for this router.	<b>ipv6 ospf neighbor</b> <i>ipv6-address</i>
Optionally, assign a cost for the specified neighbor on the interface.	<b>ipv6 ospf neighbor</b> <i>ipv6-address</i> <b>cost</b> <i>number</i>
Optionally, set the OSPFv3 priority value for the specified neighbor on the interface.	<b>ipv6 ospf neighbor</b> <i>ipv6-address</i> <b>priority</b> <i>number</i>
Optionally, set a non-broadcast neighbor polling interval.	<b>ipv6 ospf neighbor</b> <i>ipv6-address</i> <b>poll-interval</b> <i>seconds</i>
Optionally, filter outgoing link-state advertisements to an OSPFv3 neighbor on this interface.	<b>ipv6 ospf neighbor</b> <i>ipv6-address</i> <b>database-filter-all-out</b>
Optionally, set the OSPFv3 priority value for the interface.	<b>ipv6 ospf priority</b> <i>number</i>
Optionally, change the cost of sending an OSPF packet on this router interface. This setting overrides all other interface cost methods.	<b>ipv6 ospf cost</b> <i>cost</i>
Optionally, sum the interface speeds contained in the specified tracked object when calculating the OSPF interface cost.	<b>ipv6 ospf cost track</b> <i>trackobject-name</i>
Optionally, set the amount of time between retransmissions of link state advertisements (LSAs) for adjacencies that belong to the interface.	<b>ipv6 ospf retransmit-interval</b> <i>seconds</i>
Optionally, set the amount of delay before transmitting a link state update packet on an interface.	<b>ipv6 ospf transmit-delay</b> <i>seconds</i>
Optionally, enable the ignore MTU advertisement feature for the neighbor of this interface.	<b>ipv6 ospf ignore-mtu</b>
Optionally, change the number of seconds this router must wait before sending a hello packet to neighbor routers on this interface.	<b>ipv6 ospf hello-interval</b> <i>seconds</i>
Optionally, change the number of seconds this router must wait to receive a hello packet from its neighbor before determining that the neighbor is out of service.	<b>ipv6 ospf dead-interval</b> { <i>seconds</i>   <b>minimal</b> [ <b>hello-multiplier</b> <i>number</i> ]}
Optionally, configure IPsec authentication on an interface.	<b>ipv6 ospf authentication spi</b> <i>spi</i> { <b>md5</b> <i>key</i>   <b>sha1</b> <i>key</i>   <b>aescbc</b> <i>key</i> } [ <i>hex</i> ]



**Table 33-4 Configuring OSPF Optional Interface Parameters (continued)**

Task	Command(s)
Optionally, configure IPsec encrypted authentication on an interface.	<b>ipv6 ospf encryption ipsec spi <i>spi</i> {none   3des <i>key</i>   aescbc {128   192   256} <i>key</i>} [hex] auth {md5 <i>key</i>   sha1 <i>key</i>   aescbc <i>key</i>   no-auth}</b>
Optionally, disable the graceful restart helper feature on this interface.	<b>ipv6 ospf helper-disable</b>
Optionally, specify the network type that this interface is connected to.	<b>ipv6 ospf network {non-broadcast   broadcast   point-to-point   point-to-multipoint}</b>

[Table 33-5](#) describes how to display OSPFv3 configuration and statistics.

**Table 33-5 Displaying OSPFv3 Configuration and Statistics**

Task	Command(s)
Displaying OSPFv3 configuration.	<b>show ipv6 ospf</b>
Displaying OSPFv3 link-state database information.	<b>show ipv6 ospf database <i>type</i> [<i>link-state-id</i>]</b>
Displaying information about OSPFv3 internal entries to area border routers and autonomous system boundary routers.	<b>show ipv6 ospf border-routers</b>
Displaying OSPFv3 interface configuration information.	<b>show ipv6 ospf interface [<i>vlan</i> <i>vlan-id</i>]</b>
Displaying OSPFv3 neighbor information.	<b>show ipv6 ospf neighbor [<i>router-id</i>] [<i>detail</i>] [<i>vlan</i> <i>vlan-id</i>]</b>
Displaying OSPFv3 sham link information.	<b>show ipv6 ospf sham-link</b>
Displaying OSPFv3 virtual-links configuration information.	<b>show ipv6 ospf virtual-links</b>

## Intermediate System To Intermediate System (IS-IS) Configuration

This chapter provides information about configuring and monitoring Intermediate System To Intermediate System (IS-IS) on 7100-Series devices.

For information about...	Refer to page...
<a href="#">Using IS-IS in Your Network</a>	34-1
<a href="#">Implementing IS-IS</a>	34-4
<a href="#">IS-IS Configuration Overview</a>	34-4
<a href="#">Configuring IS-IS</a>	34-14
<a href="#">Terms and Definitions</a>	34-18



**Note:** IS-IS requires an advanced router license. See “[License Overview](#)” on page 9-10 for 7100-Series advanced router license details.

### Using IS-IS in Your Network

IS-IS is an interior gateway link-state routing protocol, defined in ISO 10589, and operates by reliably flooding link state information throughout a network of routers within an administrative domain.

IS-IS uses the term domain to refer to any group of routers that are administered by a single organization. For example, the network owned and operated by a single carrier would be a domain. The concept of a domain is analogous to the autonomous system in OSPF.


A domain is subdivided into areas. Areas simplify management by breaking a domain into smaller entities that are easier to manage. Routing protocols that operate within the domain are known as intra-area routing protocols (Interior Gateway Protocols (IGP) in the IP world). Routing protocols that operate between domains are known as inter-area routing protocols. IS-IS handles routing information within a domain and is therefore known as an intra-area routing protocol or IGP.

Routers that handle traffic within an area are known as Level 1 routers. Routers that handle traffic between areas are known as Level 2 routers. Routers that handle traffic both within and between areas are known as Level 1 and 2 routers. Level 1 and 2 routers run two copies of the routing algorithm.

[Figure 34-1](#) on page 34-2 displays a three area domain. Routers A, B, and C all share routes with each other. The Level 1 routers belonging to areas 1 and 2 share routes with Level 1 routers within their area. If they need to reach routers outside of their respective areas, they depend upon the Level 1 and Level 2 routers (Router A and Router B) to forward frames outside of their respective

areas. Area 3 intermediate systems beyond Router C are not displayed. Routers A and B forward frames to Router C and depend upon Router C to forward frames to other intermediate systems in Area 3.

**Figure 34-1 IS-IS Network Overview**



IS-IS is a link-state routing protocol. That is, each intermediate system (router) in a domain is represented as being in a particular state at any given time. Depending on the state of the intermediate system, different messages are expected from neighboring intermediate systems or from that intermediate system before the intermediate system can transition to the next state and ultimately exchange routing information and process its routing tables. The packets used in the IS-IS routing protocol fall into the following classes:

- Hello packets
- Link State PDUs (LSPs),

- Sequence Number PDUs (SNPs).

Hello packets are used to initialize and maintain adjacencies between neighboring intermediate systems. There are three types of IS-IS Hello packets:

- Level 1 LAN IS-IS Hello PDUs are used by Level 1 intermediate systems on broadcast LANs.
- Level 2 LAN IS-IS Hello PDUs are used by Level 2 intermediate systems on broadcast LANs.
- Point-to-Point IS-IS Hello PDUs are used on non-broadcast media, such as point-to-point links.

Link State PDUs (LSPs) contain link state information. There are two types of LSPs:

- Level 1 intermediate systems transmit Level 1 LSPs
- Level 2 intermediate systems transmit Level 1 and Level 2 LSPs

Sequence number PDUs are used to ensure that neighboring intermediate systems are aware of the most recent LSP from every other intermediate system. The sequence number PDUs serve a similar function to acknowledgement packets. There are two types of sequence number packets for both Level 1 and Level 2 intermediate systems:

- Partial sequence-number PDUs are used to request LSPs in broadcast networks and to acknowledge LSP receipt on point-to-point networks.
- A complete sequence number PDU contains a description of all LSPs in the database. A complete sequence number packet is used to ensure synchronization of the database between adjacent intermediate systems either periodically, or when an adjacency first comes up.

Intermediate systems running IS-IS establish an adjacency by passing hello packets to each other. Based on the responses to those packets, an intermediate system determines whether it can establish this adjacency and whether the adjacency is a Level 1 adjacency, a Level 2 adjacency or a Level 1 and 2 adjacency.

Because IS-IS was originally designed for OSI and then evolved to include IP routing protocols, it represents intermediate systems using OSI addressing. Each node (router or end system) in OSI is known by a unique identifier known as the Network Service Access Point (NSAP). See [“Network Layer Addresses \(NSAP\)”](#) on page 34-4 for an NSAP graphic presentation.

This address is divided into two equal parts:

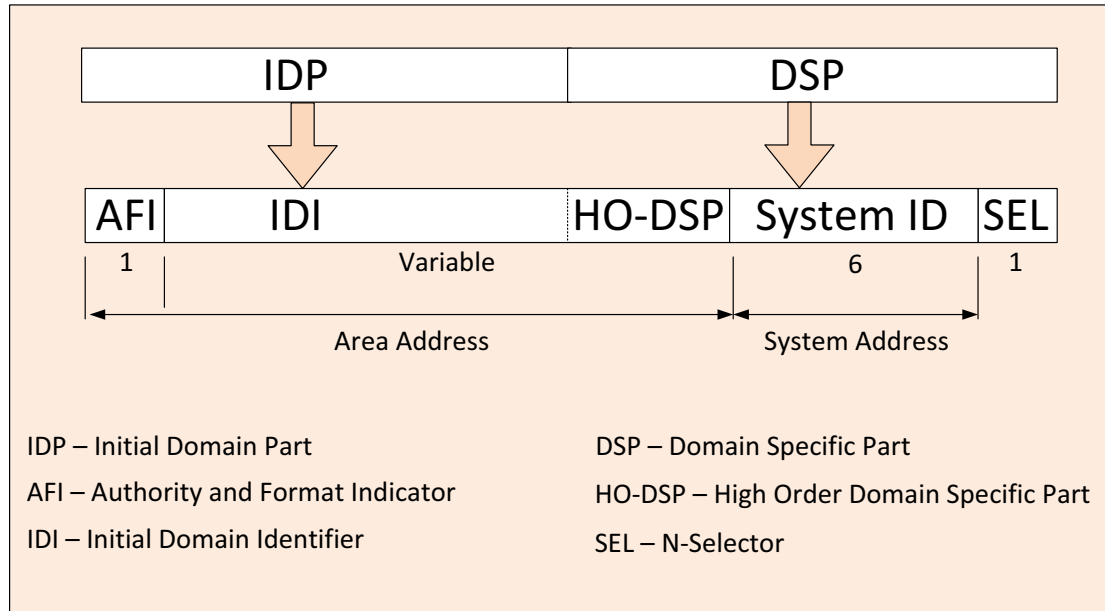
- The initial domain part (IDP)
- The domain specific part (DSP).

The initial domain part is further broken into the authority and format indicator (AFI), which specifies the format of the IDP, and the initial domain identifier (IDI). The AFI is always one octet in length.

The DSP is further broken into the High Order Domain Specific part (HO-DSP) and the system ID and selector (SEL). The system ID is always 6 octets. The SEL is one octet in length and is always 00.

The AFI, IDI and HO-DSP make up the area address which is variable in length, depending on the value of the AFI.

The network entity title (NET) of an IS-IS intermediate system is the six-byte NSAP and a single SEL byte set to 0.

**Figure 34-2 Network Layer Addresses (NSAP)**

## Implementing IS-IS

To implement IS-IS:

1. Enable IS-IS on the intermediate point.
2. Enable IS-IS on each interface that will use IS-IS routing to initiate the forming of an adjacency with the IS-IS neighbor.
3. Configure a NET for each area (up to three) the router will route to using IS-IS.
4. Optionally, modify IS-IS optional global parameters.
5. Optionally, modify IS-IS optional interface parameters.

## IS-IS Configuration Overview

For information about...	Refer to page...
<a href="#">Enabling IS-IS Globally</a>	34-5
<a href="#">Enabling IS-IS on the Interface</a>	34-5
<a href="#">Configuring a Network Entity Title (NET)</a>	34-6
<a href="#">Configuring Administrative Distance</a>	34-8
<a href="#">Configuring IS-IS Authentication</a>	34-8
<a href="#">Configuring Multiple Parallel Routes</a>	34-9
<a href="#">Enabling Route Summarization</a>	34-9
<a href="#">Configuring Route Redistribution</a>	34-11
<a href="#">Configuring IS-IS Timers</a>	34-11

For information about...	Refer to page...
<a href="#">Configuring the TLV Metric Style</a>	34-12
<a href="#">Configuring IS-IS Priority</a>	34-12
<a href="#">Configuring the IS-IS Intermediate System as Overloaded</a>	34-13

## Enabling IS-IS Globally

IS-IS must be enabled globally on the device.

Use the **router isis** command in global configuration mode to globally enable IS-IS on the intermediate system.

The following example shows how to globally enable IS-IS routing on the device:

```
7100 Chassis(rw) ->configure
7100 Chassis(rw-config) ->router isis
7100 Chassis(rw-config-isis) ->
```

The IS-IS type for the intermediate system defaults to Level 1 and 2. You can change the IS-IS type for the intermediate system to IS-IS type Level 1 or IS-IS type Level 2.

Use the **is-type** command to configure the IS-IS type for the intermediate system.

## Enabling IS-IS on the Interface

IS-IS must be enabled on each interface with an IS-IS adjacency. Not all interfaces on an intermediate system participate in IS-IS routing. Because IS-IS is a link-state routing protocol, routing information is not exchanged unless at least one adjacency is formed. In IS-IS, enabling of the IS-IS protocol on an interface begins the adjacency formation process. Therefore, no routing information is exchanged until IS-IS is enabled on at least one interface.

Use the **ip router isis** command in interface configuration mode to enable IS-IS on an IPv4 interface.

The following example shows how to enable IPv4 IS-IS routing on VLAN 100:

```
7100 Chassis(rw) ->configure
7100 Chassis(rw-config) ->interface vlan 100
7100 Chassis(rw-config-intf-vlan.0.100) ->ip router isis
7100 Chassis(rw-config-isis) ->
```

Use the **ipv6 router isis** command in interface configuration mode to enable IS-IS on an IPv6 interface.

The following example shows how to enable IPv6 IS-IS routing on VLAN 100:

```
7100 Chassis(rw) ->configure
7100 Chassis(rw-config) ->interface vlan 100
7100 Chassis(rw-config-intf-vlan.0.100) ->ipv6 router isis
7100 Chassis(rw-config-isis) ->
```

The IS-IS circuit type for the interface defaults to Level 1 and 2. You can change the IS-IS circuit type for the intermediate system to IS-IS type Level 1 or IS-IS type Level 2.

Use the **isis circuit-type** command to configure the IS-IS circuit type for the interface.

The cost of using an interface can be configured. The cost of using an interface defaults to **10**. Use the **isis metric** command in interface configuration mode to configure the interface metric.

## Configuring a Network Entity Title (NET)

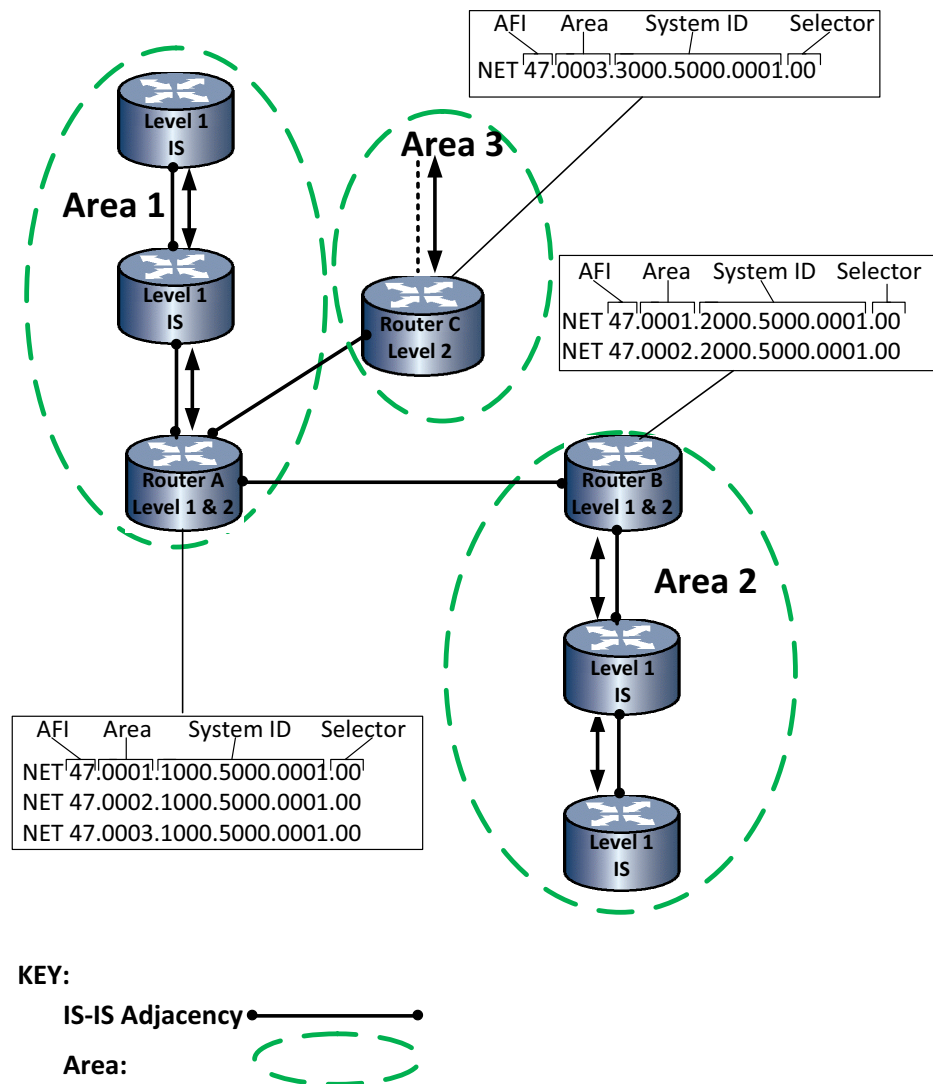
A NET is a Network Service Access Point (NSAP) address of varying length where the last byte (the NSAP-selector) is always zero. All intermediate systems within an IS-IS domain must use the same length NET. The first variable number of bytes identify the area, followed by seven fixed bytes that are divided between six bytes identifying the system ID and a single selector byte. Each intermediate system has a unique system identifier. To configure separate areas for the intermediate system, enter each area number, followed by the unique system ID for this intermediate system, followed by 00 (the NSAP-selector octet). For example: NET address **12.3333.4444.5555.6666.00** has an

- Area of **12.3333**
- System identifier of **4444.5555.6666**
- NSAP-selector of **00**

The IS-IS process does not start until at least one NET is configured.

In [Figure 34-3](#), three areas are defined: 47.0001, 47.0002, and 47.0003. Each intermediate system can belong to up to three areas. Router A is a Level 1-2 intermediate system with adjacencies in all three areas. Router B is a Level 1-2 intermediate system with adjacencies in areas 47.0001 and 47.0002. Router C is a Level 2 intermediate system with adjacencies only in area 47.0003.

Figure 34-3 IS-IS NET Configuration



Router A is configured for:

```
NET 47.0001.1000.5000.0001.00
```

```
NET 47.0002.1000.5000.0001.00
```

```
NET 47.0003.1000.5000.0001.00
```

The Level 1 adjacencies to Router A are configured with area **47.0001** and unique system IDs.

The following example configures the NETs for Router A:

```
S Chassis(rw)->configure
S Chassis(rw-config)->router isis
S Chassis(rw-config-isis)->net 47.0001.1000.5000.0001.00
S Chassis(rw-config-isis)->net 47.0002.1000.5000.0001.00
S Chassis(rw-config-isis)->net 47.0003.1000.5000.0001.00
S Chassis(rw-config-isis)->
```



Router B is configured for:

NET 47.0001.2000.5000.0001.00

NET 47.0002.2000.5000.0001.00

The Level 1 adjacencies to Router B are configured with area 47.0002 and unique system IDs.

The following example configures the NETs for Router B:

```
S Chassis(rw)->configure
S Chassis(rw-config)->router isis
S Chassis(rw-config-isis)->net 47.0001.2000.5000.0001.00
S Chassis(rw-config-isis)->net 47.0002.2000.5000.0001.00
S Chassis(rw-config-isis)->
```

Router C is configured for:

NET 47.0003.3000.5000.0001.00

Router A participates in a single-area domain. The area is area 47.0001 and the intermediate system identifier is 1000.5000.0001.

The following example configures the NETs for Router C:

```
S Chassis(rw)->configure
S Chassis(rw-config)->router isis
S Chassis(rw-config-isis)->net 47.0003.3000.5000.0001.00
S Chassis(rw-config-isis)->
```

## Configuring Administrative Distance

Administrative distance configures the preference given to a route learned from one protocol over the same route learned from another protocol. The route with the lowest administrative distance is installed in the IP routing table and propagated to neighbors.

Routes with distance values of 255 are not installed in the routing table.

Use the **distance** command in IS-IS configuration command mode to set the distance of IS-IS routes to the specified value.

IS-IS has a default administrative distance of 115.

The following example shows how to change the administrative distance for external IS-IS routes to **100**:

```
7100 Chassis(rw-config)->router isis
7100 Chassis(rw-config-isis)->distance isis external 100
```

## Configuring IS-IS Authentication

Authentication can be set both globally to authenticate between IS-IS domains and areas. Authentication can be set in interface configuration mode for authentication between intermediate systems within an area. Authentication is enabled in global configuration mode by assigning an area password. Authentication is enabled in interface configuration mode by assigning an IS-IS password.

Use the **domain-password** command to enable domain authentication at the intermediate system global level.

Use the **area-password** command to enable area authentication at the intermediate system global level.

Use the **isis password** command to enable authentication in interface configuration mode.

The authentication mode can be set at either the global or interface mode to either MD5 or text. Specify the IS-IS level of the intermediate system or interface when configuring authentication mode.

Use the **authentication-mode** command in global configuration mode to set the authentication mode at the intermediate system global level.

Use the **isis authentication-mode** command in interface configuration mode to set the authentication mode at the interface level.

A configured key-chain can be applied to IS-IS authentication in either a global or interface context.

Use the **authentication key-chain** command in IS-IS configuration mode to apply a key chain to IS-IS authentication for the intermediate system.

Use the **isis authentication key-chain** command in interface configuration mode to apply a key chain to IS-IS authentication in an interface context.

Authentication can be configured for send frames only at both the global and interface level. When configured, no authentication will be performed on received frames for the configured context.

Use the **authentication send-only** command to configure IS-IS to only include authentication on frames sent by the intermediate system.

Use the **isis authentication send-only** command in interface configuration mode to configure IS-IS to only include authentication on frames sent by the interface.

## Configuring Multiple Parallel Routes

You can set multiple parallel routes installed in the IP routing table. During packet switching, load balancing is performed among the multiple paths. A maximum of 32 paths is supported. By default, 8 paths are installed to the IP routing table.

During load balancing, a source/destination pair always uses the same path. Use the **no** option to reset the number of parallel paths to 8.

Use the **maximum-paths** command in IS-IS configuration command mode to configure the maximum number of multiple parallel routes.

## Enabling Route Summarization

Summarizing addresses reduces the number of LSPs and the size of the link state database. Multiple addresses can be summarized for a given IS-IS instance.


To summarize a unicast IPv6 address, you must be in the IPv6 unicast family address configuration mode. See “[Configuring the IPv6 Unicast Address Family](#)” on page 34-13 to enter IPv6 unicast family address configuration mode.

In IS-IS, routes are leaked from Level 1 to Level 2, on intermediate systems running both Level 1 and Level 2. The summary address command aggregates the addresses that are leaked from Level 1 and Level 2.

In [Figure 34-4](#), Router A is running Level 1 and Level 2. Router B is running Level 1 and Router C is running Level 2. Router A learns networks 192.1.1.0 and 192.1.2.0 through Level 1 IS-IS from

Router B. Without a route summary, they are leaked to Level 2 and flooded to Router C as 192.1.1.0 and 192.1.2.0.

**Figure 34-4 IS-IS Route Summarization**



The following lines are configured on Router A:

```
S Chassis(rw) ->configure
S Chassis(rw-config) ->router isis
S Chassis(rw-config-isis) ->summary-address 192.1.0.0/16
S Chassis(rw-config-isis) ->
```

Now Router A will leak summary 192.1.0.0 to Router B but will not flood routes 192.1.1.0 and 192.1.2.0. A summary is only flooded if there is at least one Level 1 route that falls into the configured summary address range.

## Configuring Route Redistribution

You can redistribute routes into IS-IS for route types and protocols: BGP, connected, OSPF, RIP, Static and blackhole. You can also redistribute IS-IS Level 2 routes into Level 1.

A route map can be applied to a redistribution configuration to filter the routes that will be redistributed for the specified route type or protocol. You can also specify a protocol metric value for BGP, connected, RIP, or static redistribution routes.

You must be in IS-IS IPv6 unicast address family configuration mode to configure redistribution on an IPv6 IS-IS intermediate system instance. See [“Configuring the IPv6 Unicast Address Family”](#) on page 34-13 for details on entering IS-IS IPv6 unicast address family configuration mode.

If you do not specify a distribute list when redistributing IS-IS level 2 routes into IS-IS level 1, all layer 2 addresses are redistributed into layer 1.

Use the **redistribute** command in IS-IS configuration or IS-IS IPv6 unicast address family configuration mode to specify the protocol or route type to redistribute into IS-IS.

## Configuring IS-IS Timers

There are three global level and four interface level IS-IS timers. See [Table 34-1](#) for IS-IS timer details.

**Table 34-1 IS-IS Timers**

Timer	Description	Command
LSP Generation Interval	The minimum interval between the generation of LSPs, configured in the IS-IS configuration mode.	<b>isp-gen-interval</b>
Maximum LSP Lifetime	The maximum time the LSPs persist without being refreshed.	<b>max-isp-lifetime</b>
SFP Interval	The minimum amount of time between Shortest Path First (SFP) processing on an IS-IS instance. When a topology change occurs the SPF calculation is run. The SPF calculation is not run when external routes change.  The SPF calculation is CPU intensive. For a network with a large area and frequent topology changes you may want to increase the minimum time between SPF calculations. Increasing the SPF interval reduces the processor load, but potentially slows the rate of convergence.	<b>sfp-interval</b>
Complete Sequence Number PDU Interval	The IS-IS Complete Sequence Number PDU (CSNP) interval for the interface. Designated Routers (DRs) send out CSNP packets on the interface to maintain database synchronization.	<b>isis csnp-interval</b>

**Table 34-1 IS-IS Timers (continued)**

Timer	Description	Command
Hello Interval	The IS-IS Hello PDU interval for the interface. The advertised holdtime in the hello packet is set to three times the hello interval seconds. Topological changes are detected faster with a smaller hello interval, but there is more routing traffic.  The number of hello packets a neighbor must miss before the intermediate system declares the adjacency down on the interface is the holdtime multiplier which can be changed using the <b>isis hello-multiplier</b> command.	<b>isis hello-interval</b>
LSP transmission interval	The minimum interval between the transmission of Link-State Packets (LSP)s on the interface.	<b>isis lsp-throttle</b>
LSP retransmission interval	The minimum interval between the retransmission of the same LSP. The retransmit interval should be greater than the expected round-trip delay between any two intermediate systems on the attached network. Retransmissions occur when LSPs are dropped. Higher retransmission values have little effect on reconvergence. The more neighbors intermediate systems have, and the more paths over which LSPs can be flooded, the higher this value can be made.	<b>isis retransmission-interval</b>

## Configuring the TLV Metric Style

By default the IS-IS metric style defaults to both wide and narrow. The narrow metric is a 6-bit metric as defined in ISO 10589. The wide metric is a 4-byte metric and configures an intermediate system to generate and accept type, length, and value (TLV) object 135 for IP addresses. The TLV metric style can be configured for either both wide and narrow or wide.

Use the **metric-style** command in IS-IS configuration mode to configure the TLV metric style for the IS-IS intermediate system.

## Configuring IS-IS Priority

The priority is used to determine the designated router. The intermediate system with the highest priority becomes the designated router. IS-IS does not support the concept of a backup designated router. Setting the priority to 0 does not prevent this system from becoming the designated router. If priorities are equal, the interface with the highest MAC address breaks the tie. Priority defaults to **64**.

The following example shows how to configure VLAN **100** for a priority of **80**:

```
7100 Chassis(rw) ->configure
7100 Chassis(rw-config) ->interface vlan 100
7100 Chassis(rw-config-intf-vlan.0.100) ->isis priority 80
7100 Chassis(rw-config-intf-vlan.0.100) ->
```

## Configuring the IS-IS Intermediate System as Overloaded

There are some circumstances in which it is advantageous to have an IS-IS intermediate system not fully participate in forwarding traffic. For example:

- During startup this intermediate system may be temporarily too busy with the tasks associated with convergence to forward traffic.
- The intermediate system is in a test network that has connections to a production network. The overload bit prevents traffic from moving between the two networks.

You can configure the intermediate system to not forward traffic by enabling the overload bit.

When enabled, the overload bit instructs other intermediate systems not to use this intermediate system as an intermediate hop in their SPF computations. No paths through this intermediate system are visible to other intermediate systems in the domain. IP and CLNS prefixes directly connected to this intermediate system are reachable.

Use the **set-overload-bit** command in IS-IS configuration mode to enable overload on this intermediate system.

## Configuring the IPv6 Unicast Address Family

By default, IS-IS configuration applies to both IPv4 and IPv6 routes. Use this command to configure IPv6 specific configuration on the device. IPv6 unicast specific address family configuration currently supports:

- Administrative distance – See “[Configuring Administrative Distance](#)” on page 34-8
- Redistribution of routes from other protocols into IS-IS – See “[Configuring Route Redistribution](#)” on page 34-11
- Route summarization– See “[Enabling Route Summarization](#)” on page 34-9

The following example shows how to enter IPv6 unicast IS-IS address family configuration mode and set the IPv6 unicast IS-IS administrative distance to 100:

```
7100 Chassis(rw) ->configure
7100 Chassis(rw-config)->router isis
7100 Chassis(rw-config-isis)->address-family ipv6 unicast
7100 Chassis(rw-config-isis-af)->distance 100
7100 Chassis(rw-config-isis-af)->
```

## Graceful Restart

IS-IS graceful restart, sometimes referred to as non-stop forwarding, provides for an IS-IS router to remain on the forwarding path during a restart of its IS-IS software. Graceful-restart has four elements to its configuration: enabling, helping a peer, database re-synchronization interval, and restart interval.

Enabling graceful restart instructs the firmware to perform a graceful restart, rather than a standard IS-IS restart. Restart is only initiated by a fail-over. Whether the failover is intentional or not, the failed router protocol is restarted on another module, and upon startup, IS-IS informs its neighbors using existing link aggregation groups. Use the **graceful-restart enable** command to enable the graceful restart ability on this router.

When the helper peer is informed that a graceful restart is taking place, it sends the restarting router its database and prevents the rest of the network from being informed there is an issue with the restarting router. The helper also monitors the network for topology changes. If no changes occur, the helper router continues to advertise as though no restart was occurring. If a topology

change does occur, graceful restart is terminated on the restarting router and a standard restart occurs. Helper mode is enabled by default, but can be disabled on an IS-IS router using the **no graceful-restart enable-help-peer** command in IS-IS router command mode.

The length of time to allow database synchronization during a graceful restart can be configured using the **graceful-restart restart-sync-interval** command in IS-IS router configuration mode.

A restart interval provides for a maximum time in seconds after which the graceful restart will terminate should it not complete or terminate for other reasons within the interval. Use the **graceful-restart restart-interval** command to change the restart interval setting.

View the router OSPF section of the **show running-config** display to verify any non-default graceful restart settings.

## Configuring IS-IS

This section provides a table of feature default values and a procedure for configuring a feature system.

[Table 34-2](#) lists IS-IS default values.

**Table 34-2 IS-IS Parameters**

Parameter	Description	Default Value
authentication mode	The IS-IS MD5 authentication mode that provides a cryptographic hash MD5 digest to each IS-IS PDU.	MD5
distance	The administrative distance for IS-IS routes.	115
graceful restart state	Specifies whether graceful restart is enabled or disabled on a router.	disabled
graceful restart help peer	A function on the graceful restarting router peer that specifies whether the peer will help in the graceful restarting process.	enabled
graceful restart adjacency interval	Specifies the length of time graceful restart waits for the adjacency to form.	10 seconds
graceful restart interval	Specifies the length of time a graceful restart is allowed to continue without completing before graceful restart is terminated.	65535 seconds
graceful restart sync interval	Specifies the length of time graceful restart will wait for the IS-IS database to synchronize.	60 seconds
ignore LSP errors	Specifies that IS-IS ignores link state packet checksum errors on the device.	disabled
LSP buffer size	Specifies the LSP buffer size, based upon the maximum size of LSPs originated by this IS-IS routing instance.	1492 bytes
LSP generation interval	The minimum interval between generation of LSPs.	1 second
maximum LSP lifetime	The maximum time an LSP can persist without being refreshed.	1200 seconds
maximum paths	The maximum number of parallel routes to be installed into the routing table for this device.	8
metric style	The TLV metric style for this IS-IS instance.	both (narrow and wide)

**Table 34-2 IS-IS Parameters (continued)**

Parameter	Description	Default Value
overload bit	When set, instructs this intermediate system to tell other intermediate systems not to use it as an intermediate hop in the SPF calculation.	not set
SPF interval	The minimum amount of time between Shortest Path First (SPF) processing for this IS-IS instance.	33 milli-seconds
CSNP interval	The interval between sending sequence number PDUs (CSNP) on the interface.	10 seconds
hello interval	The minimum amount of time between sending hello PDUs on the interface.	10 seconds
hello multiplier	The number of hello packets a neighbor must miss before the intermediate system declares the adjacency down for the interface.	3
LSP MTU	The maximum PDU size for PDUs on the interface.	1490 bytes
LSP throttle	The minimum interval between the transmission of Link-State-Packets on the interface.	30 milli-seconds
metric	The cost of using the interface.	10
priority	The priority used to determine which intermediate system on a LAN is the designated router.	64
retransmit interval	The minimum amount of time between the retransmission of the same LSP on an interface.	5 seconds

[Procedure 34-1](#) describes feature configuration on the Extreme Networks 7100-Series devices. All commands used to configure IS-IS are entered in the IS-IS configuration mode after enabling IS-IS using the **router isis** command.

#### **Procedure 34-1 Configuring Global IS-IS**

Step	Task	Command(s)
1.	In global configuration mode, enable IS-IS globally on the device and enter IS-IS intermediate system configuration mode.	<b>router isis</b>
2.	Specify the Network Access Point address for this IS-IS instance in the format <code>xx.xxxx.(...)xxxx.xxxx.xxxx.00</code> .	<b>net net</b>
3.	Optionally configure the IS-IS type for this IS-IS instance.	<b>is-type {level-1   level-1-2   level-2}</b>
4.	Specify the IS-IS domain password for this device.	<b>domain-password password</b>
5.	Specify the area password.	<b>area-password password</b>
6.	Optionally specify an IS-IS authentication key chain for the device.	<b>authentication key-chain keychain [level-1   level-1-2   level-2]</b>



**Procedure 34-1 Configuring Global IS-IS (continued)**

Step	Task	Command(s)
7.	Optionally set the authentication mode.	<b>authentication mode</b> {md5   text} [level-1   level-1-2   level-2]
8.	Optionally configure IS-IS authentication only on sent IS-IS frames.	<b>authentication send-only</b> [level-1   level-1-2   level-2]
9.	In IS-IS configuration mode, optionally modify the IS-IS administrative distance for this IS-IS intermediate system.	<b>distance</b> [isis {external   internal}] <i>weight</i>
10.	Optionally configure a dynamic hostname.	<b>hostname dynamic</b> <i>hostname</i>
11.	Optionally enable IS-IS to ignore link state packet checksum errors on the device.	<b>ignore-lsp-errors</b>
12.	Optionally configure the LSP buffer size based upon the specified maximum size of the LSP's originated by the IS-IS routing instance.	<b>lsp-buf-size</b> <i>size</i> [level-1   level-1-2   level-2]
13.	Optionally configure the minimum interval between the generation of LSPs on the device.	<b>lsp-gen-interval</b> <i>interval</i> [level-1   level-1-2   level-2]
14.	Optionally configure the maximum time that LSPs persist without being refreshed.	<b>max-lsp-lifetime</b> <i>lifetime</i>
15.	Optionally configure the maximum number of parallel routes to be installed into the routing table for this device.	<b>maximum-paths</b> <i>num</i>
16.	Optionally configure the TLV metric style for this IS-IS instance.	<b>metric-style</b> {wide   both}
17.	In IS-IS configuration mode, optionally configure the redistribution of routing protocols into IS-IS.	<b>redistribute</b> {bgp   connected   ospf <i>process_id</i>   rip   static   blackhole   isis level-2 into level-1 [distribute-list <i>access-list</i> ]} [route-map <i>name</i> ] [metric <i>metric-value</i> ]
18.	Optionally configure the intermediate system to signal other intermediate systems not to use this intermediate system as an intermediate hop in their SPF calculations.	<b>set-overload-bit</b> [level-1   level-1-2   level-2]
19.	Optionally configure the minimum amount of time between Shortest Path First (SPF) processing on an IS-IS instance.	<b>spf-interval</b> <i>interval</i>
20.	Optionally create an aggregate IS-IS address for summarization of routes.	<b>summary-address</b> <i>ip-address/length</i>
21.	Optionally enable IS-IS graceful restart for this router.	<b>graceful-restart enable</b>
22.	Optionally set the length of time graceful restart will allow for the forming of adjacencies.	<b>graceful-restart restart-adj-interval</b> <i>interval</i>
23.	Optionally set the length of time graceful restart will attempt to complete a restart before terminating.	<b>graceful-restart restart-interval</b> <i>interval</i>
24.	Optionally specify the amount of time graceful restart will allow for the synchronization of the database.	<b>graceful-restart restart-sync-interval</b> {level-1   level-1-2   level-2} <i>interval</i>

Table 34-3 describes IS-IS Interface configuration on the Extreme Networks 7100-Series devices. All IS-IS interface configuration commands are entered in interface configuration mode. All IS-IS interface configuration is optional.

**Table 34-3 Configuring IS-IS on the Interface**

Task	Command(s)
If IPv6 routing will take place on the interface, enable IPv6 IS-IS routing on the interface. IPv4 IS-IS routing is enabled globally.	<b>ipv6 router isis</b>
Configure an IS-IS authentication key chain on an interface	<b>isis authentication key-chain <i>keychain</i> [level-1   level-2]</b>
Configure the IS-IS authentication mode on an interface.	<b>isis authentication mode {md5   text} [level-1   level-2]</b>
Configure IS-IS authentication only on sent IS-IS frames on an interface.	<b>isis authentication send-only [level-1   level-2]</b>
Configure the IS-IS type on an interface	<b>isis circuit-type {level-1   level-1-2   level-2}</b>
Configure the IS-IS complete sequence number PDU (CSNP) interval for the interface.	<b>isis csnp-interval <i>seconds</i> [level-1   level-1-2   level-2]</b>
Configure the IS-IS Hello Protocol Data Units interval for the interface.	<b>isis hello-interval {<i>seconds</i>   minimal} [level-1   level-1-2   level-2]</b>
Configure the number of hello packets a neighbor must miss before the intermediate system declares the adjacency down for the interface.	<b>isis hello-multiplier <i>multiplier</i> [level-1   level-1-2   level-2]</b>
Configure IS-IS hello padding on an interface.	<b>isis hello-padding</b>
Configure the maximum PDU size for PDUs on the interface.	<b>isis lsp-mtu <i>size</i></b>
Configure the minimum interval between the transmission of Link-State Packets (LSP)s.	<b>isis lsp-throttle <i>interval</i></b>
Configure the cost of using the interface.	<b>isis metric <i>cost</i> [level-1   level-1-2   level-2]</b>
Configure a two device network that uses broadcast media and IS-IS to function as a point-to-point link.	<b>isis network-point-to-point</b>
Configure the suppression of IS-IS packets from being transmitted by the interface and received packets from being processed by the interface.	<b>isis passive-interface</b>
Configure an authentication password for the interface.	<b>isis password <i>password</i> [level-1   level-2]</b>
Configure the priority used to determine which intermediate system on a LAN is the designated router.	<b>isis priority <i>priority</i> [level-1   level-1-2   level-2]</b>
Configure the minimum interval between retransmissions of the same LSP.	<b>isis retransmit-interval <i>interval</i></b>

[Procedure 34-2](#) describes IS-IS IPv6 unicast address family configuration on the Extreme Networks 7100-Series devices.

### Procedure 34-2 Configuring IS-IS IPv6 Unicast Address Family

Step	Task	Command(s)
1.	From global configuration mode, optionally enter the IPv6 unicast IS-IS address family configuration mode to configure administrative distance, redistribution of routing protocols into IS-IS, and address summarization for IPv6 unicast.	<b>address-family ipv6 unicast</b>
2.	In IPv6 unicast address family configuration mode, optionally modify the IS-IS administrative distance for this IS-IS intermediate system.	<b>distance [isis {external   internal}] weight</b>
3.	In IPv6 unicast address mode, optionally configure the redistribution of routing protocols into IS-IS.	<b>redistribute {bgp   connected   ospf process_id  rip   static   blackhole   isis level-2 into level-1 [distribute-list access-list]} [route-map name] [metric metric-value]</b>
4.	In IPv6 unicast address family configuration mode, optionally create an aggregate IS-IS address for summarization of routes.	<b>summary-address ip-address/length</b>

[Table 34-4](#) describes how to display IS-IS information on the Extreme Networks 7100-Series devices. All IS-IS display commands can be entered from any command mode.

### Table 34-4 Displaying IS-IS Information

Task	Command(s)
Display IS-IS database information for the intermediate system.	<b>show isis database [Isp /sp]   [[level-1]   [level-2]   [detail]   [verbose]</b>
Display the hostname per LSP ID.	<b>show isis hostname</b>
Display the frequency and reason for LSP changes on an interface.	<b>show isis lsp-log</b>
Display IS-IS intermediate system neighbors.	<b>show isis neighbors</b>
Display the IS-IS topology.	<b>show isis topology</b>

Refer to the *Extreme Networks 7100-Series CLI Reference* for more information about each command.

## Terms and Definitions

[Table 34-5](#) lists terms and definitions used in this feature configuration discussion.

### Table 34-5 Feature Configuration Terms and Definitions

Term	Definition
adjacency	A direct connection between IS-IS routers.
area	An area is a logical segmenting of an IS-IS domain for purposes of simplifying network management.
domain	An IS-IS network administered by a single organization.

**Table 34-5 Feature Configuration Terms and Definitions (continued)**

<b>Term</b>	<b>Definition</b>
hello packets	An IS-IS packet type used to initialize and maintain adjacency between neighboring intermediate systems.
intermediate system	A router in an IS-IS context.
Intermediate System to Intermediate System (IS-IS)	IS-IS is an interior gateway link-state routing protocol, defined in ISO 10589, operating by reliably flooding link state information throughout a network of routers within an administrative domain.
link state PDU	An IS-IS packet type used to exchange a router's link state information with other routers in the IS-IS network.
Network Entity Title (NET)	The portion of the NSAP that contains the system ID and the selector field.
Network Service Access Point (NSAP)	An addressing scheme to identify IS-IS intermediate points and the area they belong to.
router level	One of three possible levels (L1, L2, and L1 and L2) assigned to a router to determine whether route exchange is limited to an area or can be between areas.
sequence number PDUs	An IS-IS packet type used to ensure that neighboring intermediate systems are aware of the most recent LSP from every other intermediate system.



## Border Gateway Protocol (BGP) Configuration

This chapter provides the following information about configuring and monitoring BGP on Extreme Networks 7100-Series devices:

For information about...	Refer to page...
<a href="#">Using BGP in Your Network</a>	35-1
<a href="#">Implementing BGP</a>	35-4
<a href="#">BGP Overview</a>	35-5
<a href="#">Configuring BGP</a>	35-22
<a href="#">Terms and Definitions</a>	35-51



**Note:** BGP requires an advanced router license. See “[License Overview](#)” on page 9-10 for 7100-Series advanced router license details.

### Using BGP in Your Network

The Border Gateway Protocol (BGP), documented in RFC 4271, is the standard protocol for routing between administrative domains. BGP refers to an administrative domain as an Autonomous System (AS). BGP is an exterior gateway routing protocol (EGP). BGP exchanges routing information among neighboring routers in different autonomous systems. An autonomous system is a set of routers under a single administration. AS numbers supported are 1 to 4294967295.

An AS typically uses a single Interior Gateway Protocol (IGP), such as OSPF, to propagate routing information among its routers.

A BGP system establishes sessions with neighboring routers, or peers, and maintains a database of network reachability information that it exchanges with its neighbors via update messages. BGP uses the Transmission Control Protocol (TCP) and port 179 for establishing connections.

#### Path Attributes

BGP routing updates include the complete route to each destination, as well as other information related to the route. Route information is included in the path attributes. BGP uses path attributes to provide more information about each route. Path attributes can also be used to distinguish between groups of routes to determine administrative preferences, allowing greater flexibility in determining route preference to achieve a variety of administrative ends. Supported BGP attributes include IP next hop, Multi-Exit Discriminator (MED), and local preference. BGP also uses path attributes to maintain the AS path.

The AS path is a path attribute that provides a list of the AS numbers the route traverses. The AS path is used for loop detection. Its length is used as a route selection criteria in the event the same

prefix is learned from multiple peers. BGP uses the AS path and the path attributes to determine the network topology. This, in turn, enables BGP to detect and eliminate routing loops and to make routing policy decisions.

Refer to “[Using AS-Path Regular Expressions](#)” on page 35-7 for information about using regular expressions when configuring AS path preference in route-maps.

### **Peers and Peer Groups**

A peer is the BGP router’s next hop neighbor. Peers can be organized into peer groups. A peer group is a group of neighbors that share the same attributes. You assign neighbors to the peer group, and create attributes for the group. Peer groups enable you to reduce the time and effort needed for configuration. You can override a peer group’s policy by configuring unique policies for a specific peer group member.

### **BGP Sessions**

BGP supports two basic types of sessions between neighbors: internal (sometimes referred to as IBGP) and external (EBGP). Internal sessions are run between routers in the same autonomous system. External sessions run between routers in different autonomous systems. When a router routes to an external peer, the local AS number is prepended to the AS path. This means that routes received from an external peer are guaranteed to have the AS number of that peer at the start of the path. In general, routes received from an internal neighbor will not have the local AS number prepended to the AS path. Those routes will have the same AS path that the route had when the first internal neighbor received the route from an external peer. Routes with no AS numbers in the path may be legitimately received from internal neighbors. BGP considers these routes internal to the receiver’s own AS.

External BGP sessions may or may not include the Multi-Exit Discriminator (MED) among its path attributes. BGP uses MED to break ties between routes with equal preference from the same neighboring AS.

Internal BGP sessions carry the local preference attribute. The larger the local preference value, the greater the route is preferred within an AS. Internal sessions can optionally include the MED, carried in from external sessions.

### **Routes**

A route consists of a prefix, a prefix length, and a set of information indicating policies and preference to reach the destinations indicated by the prefix. A prefix is made up of a dotted decimal formatted network identifier that includes a length that specifies the number of significant bits in the network. The route prefix is contained in the Network Layer Reachability Information (NLRI) and the BGP next hop path attribute determines where packets matching the prefix should be forwarded. The BGP next hop may be non-directly connected. In this case, for the route to be installed in the routing table, the router must have a route to the BGP Next Hop.

You can redistribute routing information between BGP and another protocol, and use route-maps to control the route updates.

### **Routing Policy**

Routing policies can be used to filter routes both on an import and export basis, based upon its IP prefix, community (RFC 1997), extended community (RFC 4260), AS path, source IP address, and IP next hop. Routing policy is configured in a route-map, which is then applied to the route.

### **Confederations and Route Reflectors**

Confederations enable you to divide a large AS into several smaller ASs, or to create an AS out of members of multiple ASs. Confederations are still fully meshed but require only a single connection to other ASs, reducing the number of peering relationships. From the outside, the confederation of smaller ASs looks like a single AS. Confederations are defined in RFC 3065

Route reflection enables you to configure a BGP router to advertise the routes learned from clients to other clients. This eliminates the full mesh requirement. You can configure one or more routers in the AS to be reflectors. The other routers are configured as clients. Route reflection is defined in RFC 4456.

## BGP Sub-Features

Supported BGP sub-features include:

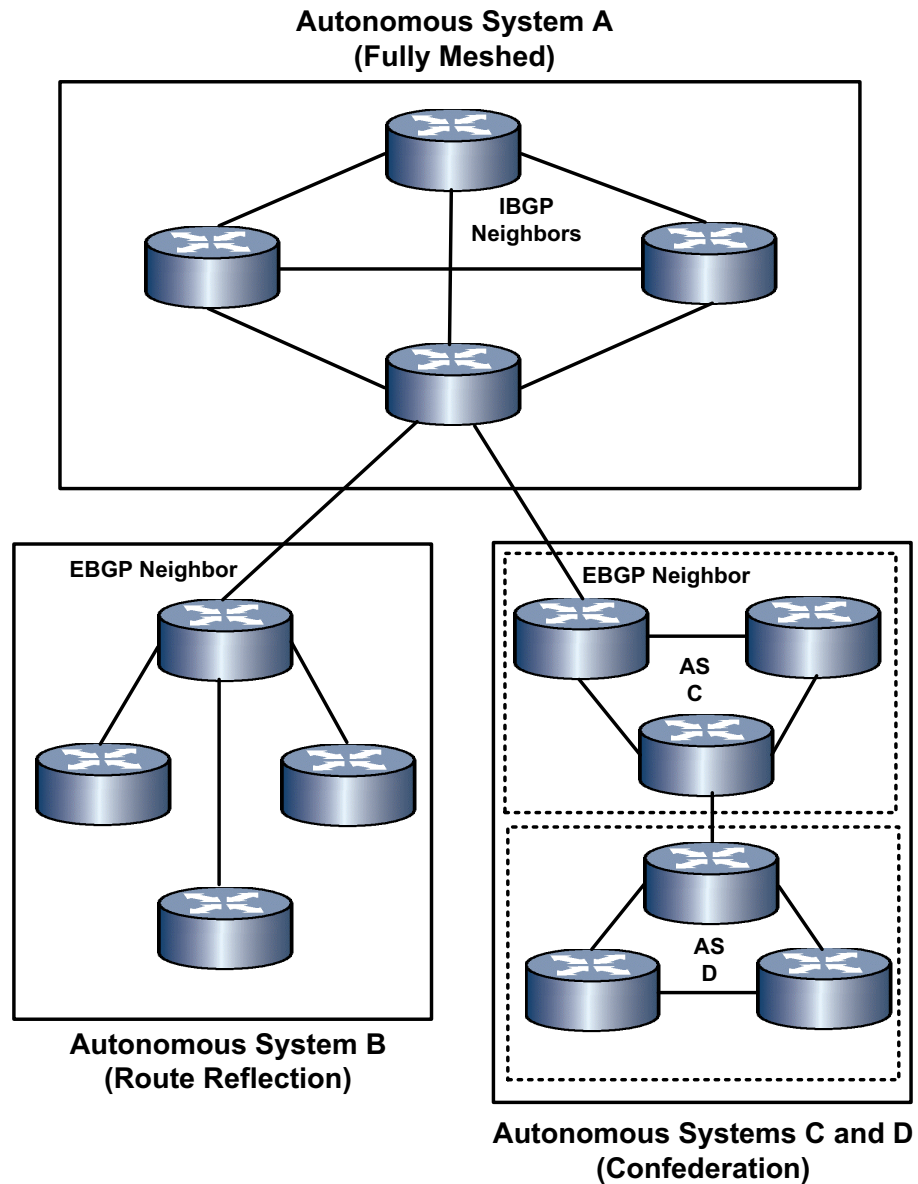
- Graceful restart – Provides for the continued processing of the data-forwarding plane of a router should the control plane fail (RFC 4724)
- Outbound Route Filtering – Allows a BGP speaker to send to its BGP peer a set of Outbound Route Filters (ORFs), which the peer applies in addition to its locally configured outbound filters (if any), to constrain its outbound routing updates to the speaker (RFC 5291)
- Route Refresh – Allows for the dynamic exchange of route refresh requests between BGP speakers and subsequent re-advertisement of the respective Adj-RIB-Out (RFC 2918)
- Route-Flap Dampening – Treats routes that are being announced and withdrawn at a rapid rate as unreachable, based upon a route penalty for each route withdrawal, and reachable again, based upon a configurable decay over time of that route penalty (RFC 2439)
- Multiprotocol BGP Extensions – Enable BGP to carry routing information for multiple Network Layer protocols such as IPv6 and IPX (RFC 2858)
- 4-Octet AS numbers – Allows for the encoding of 4-octet AS numbers (RFC 4893)
- TCP/MD5 Authentication – Enhances BGP security by defining a TCP option for carrying an MD5 digest in a TCP segment that acts like a signature for that segment, incorporating information known only to the connection end points (RFC 2385)
- Conditional Advertisement – Provides for the sending of BGP announcements, in addition to normal announcements, when a route specified in the configured advertise map does not exist in the configured non-exist map
- Aggregation – Provides for the aggregating of one or more specific routes into a single aggregate route, if a more specific route of the aggregate route exists in the BGP routing table.
- Soft Reconfiguration – Speeds up the route installation process when an inbound policy change occurs by keeping a local copy of the routes for the specified peer or group

Figure 35-1 shows a sample BGP topology with four autonomous systems:

- Autonomous system **A** displays a standard fully meshed AS
- Autonomous system **B** displays a route reflected topology
- Autonomous system **C** and **D** displays a confederation topology with two confederations

### Figure 35-1 BGP Topology





## Implementing BGP

Before configuring BGP on the routers in your network, map out the network BGP topology including autonomous systems (full mesh, route reflected, and confederation), member routers, router peers, peer policy (route-maps)

### Required steps to implement BGP in your network:

- Configure each router specifying the autonomous system the router belongs to
- Configure each router as part of a full mesh, route reflection, or confederation topology
- Configure all IBGP and EBGP neighbors for the router including all optional neighbor parameters specified in "[Configuring BGP Neighbor Parameters](#)" on page 35-34

**BGP parameters with default values that can be modified:**

- Optionally modify the route MED value using an applied route-map and optional MED behaviors using the appropriate BGP commands
- Optionally modify the local preference of advertised routes for the router
- Optionally modify the BGP route selection priority (distance) compared to other protocols for the router
- Optionally modify maximum allowed EBGP and IBGP ECMP routes for the router

**BGP features that can be configured on the router:**

- Optionally configure aggregate addresses
- Optionally configure soft reset on the router by configuring soft reconfiguration for the neighbor or automatic router refresh for the router
- Optionally configure flap dampening on the router
- Optionally enable graceful restart on the router
- Optionally configure outbound route filtering for the router
- Optionally configure BGP route-maps and apply them to configured neighbors and route redistribution
- Optionally configure Syslog and trap behavior for changes in peer state

## BGP Overview

For information about...	Refer to page...
<a href="#">Injecting Routes Into BGP</a>	<a href="#">35-5</a>
<a href="#">Using AS-Path Regular Expressions</a>	<a href="#">35-7</a>
<a href="#">Route Selection Preference</a>	<a href="#">35-8</a>
<a href="#">Multi-Exit Discriminator (MED)</a>	<a href="#">35-8</a>
<a href="#">Route Aggregation</a>	<a href="#">35-9</a>
<a href="#">Source IP Address Update to the Peer</a>	<a href="#">35-10</a>
<a href="#">Scalability and the Peer Full Mesh Requirement</a>	<a href="#">35-11</a>
<a href="#">Outbound Route Filtering (ORF)</a>	<a href="#">35-13</a>
<a href="#">Conditional Advertisement</a>	<a href="#">35-13</a>
<a href="#">BGP Soft Reset</a>	<a href="#">35-14</a>
<a href="#">Community and Extended Community Attributes</a>	<a href="#">35-15</a>
<a href="#">Route Flap Dampening</a>	<a href="#">35-18</a>
<a href="#">Graceful Restart</a>	<a href="#">35-21</a>

### Injecting Routes Into BGP

Routes can be injected into BGP using either redistribution for connected, static, RIP, or OSPF routes or by specifying a network prefix and length using the **network** command. The route or

routes are injected into BGP as long as they are valid routes with resolvable next hops. Once injected into BGP, these routes can be advertised to BGP peers.

## Using Redistribution

With redistribution, the user specifies the source protocol in BGP router or address family configuration mode. Redistribution can be configured for all routes of the specified type or routes can be filtered using a route-map.

Redistribution entries are created with a specified source and destination protocol to allow redistribution from the source to the destination protocol. The user can also configure a route-map to specify a set of matching prefixes as well as to set route attributes on matching routes.

In the 7100-Series implementation for the redistribution route-map, matching is performed only on an IP prefix as specified in an access-list. The redistribution command line allows for setting MED, local-preference, AS-path limit, and origin to matching routes. Filtering on AS-path regular expressions is supported; see [“Using AS-Path Regular Expressions”](#) on page 35-7.

BGP route-maps support setting the AS, AS-path limit, community, a number of extended community values, local preference, MED, IP next hop, origin, ORF-association local, weight, and flap table. See [Chapter 39, Route-Map Manager Configuration](#) for BGP route-map configuration details.

Use the **redistribute connected** command, optionally specifying a route-map, AS-path limit, origin, MED, and local preference for the route, to inject all or filtered connected routes into BGP.

In the following example BGP is configured to redistribute connected routes that match the contents of the **connRoute** route-map:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->redistribute connected route-map connRoute
```

Use the **redistribute rip** command, optionally specifying a route-map, AS-path limit, origin, MED, and local preference for the route, to inject all or filtered RIP routes into BGP.

In the following example BGP is configured to redistribute all RIP routes with the local preference set for **100**.

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->redistribute rip local-pref 100
```

Use the **redistribute static** command, optionally specifying a route-map, AS-path limit, origin, MED, and local preference for the route, to inject all or filtered static routes into BGP.

In the following example BGP is configured to redistribute all static routes with the local preference set for 100.

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->redistribute static local-pref 100
```

Use the **redistribute ospf** command, optionally specifying the route-map, AS-path limit, origin, MED, and local preference attributes for the route, to inject all or filtered OSPF routes into BGP.

In the following example BGP is configured to redistribute OSPF routes that match the contents of the **OSPFroutes** route-map.

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->redistribute ospf route-map OSPFroutes
```

## Using the Network Command

With the **network** command, the user explicitly specifies the Network IP prefix to be injected into BGP. The route will be injected into BGP for advertisement to BGP peers as long as the local router

has a route to the prefix with a reachable next-hop. The **network** command supports the injection of the default route (0.0.0.0/0) into BGP, if the route is present in the routing table.

Use the **network** command, specifying the network prefix and length and optionally specifying the route-map, AS-path limit, origin, MED, and local preference attributes for the route, to inject a route into BGP.

The following example imports the network 10.1.0.0 with a mask of 255.255.255.0 into BGP. Additionally, this network range will be advertised to other peers.

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp router-id 159.1.1.9
7100 Chassis(su-config-bgp)->network 10.1.0.0/24
```

The following example imports the prefix 2001::/64 into BGP. This network will be advertised based upon the **routes1** route-map contents with origin set to IGP.

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(config-router-bgp)-> bgp router-id 1.2.3.4
7100 Chassis(config-router-bgp)-> network 2001::/64 route-map routes1 origin 0
```

## Using AS-Path Regular Expressions

The BGP AS path attribute includes a list of autonomous systems that routing information has passed through to get to a specified router and an indicator of the origin of this route. Each autonomous system through which a route passes prepends its AS number to the beginning of the AS path. The AS path is used to prevent routing loops in BGP.

This routing information can be used to prefer one path to a destination network over another. The route-map supports the configuration of AS path preference when importing and exporting routes. The matching of an AS path packet attribute is configured in a route-map using the **match as-path** command.

The **match as-path** command supports the following regular expressions:

**Table 35-1 AS-Path Regular Expressions**

Character	Description	Example
<b>^</b>	Start of string.	<b>^200</b> matches any string starting with "200".
<b>\$</b>	End of string.	<b>200\$</b> matches any string ending with "200".
<b>.</b>	Matches any character.	<b>2.0</b> match "200", "210", "220"
<b>*</b>	Matches the preceding element zero or more times.	<b>22*</b> matches "2", "22", "222"
<b>.*</b>	Matches any character any number of times.	This is a match all.
<b>[]</b>	Matches a single character inside the brackets.	<b>[d]</b> matches the character "d".
<b>[-]</b>	Matches a range.	<b>[0-9]</b> matches any number from 0 to 9.
<b>()</b>	Specifies a subexpression.	<b>(200:500)</b> is treated as a single entity.
<b>[^]</b>	Matches any single character not specified in the brackets.	<b>[^er]</b> matches all characters except for "e" and "r".
<b>?</b>	Matches the preceding exactly.	<b>200?</b> matches "200" only.

**Table 35-1 AS-Path Regular Expressions (continued)**

Character	Description	Example
+	Matches the preceding element one or more times.	200? matches "200", "200200", "200200200".
	Matches either the expression before or the expression after the operator (the choice operator).	100   250 matches either "100" or "250"



**Note:** Regular expressions are also supported by the BGP community and extended community attributes.

This example shows how to match a packet AS path attribute that starts with AS number **20313** and with the next AS number ending with **13**:

```
7100 Chassis(su)->configure
7100 Chassis(su-config)->route-map bgp bgprm1 permit
7100 Chassis(su-config-route-map-bgp)->match as-path ^20313.*13$
7100 Chassis(su-config-route-map-bgp)->show route-map bgprm1
route-map bgp bgprm1 permit 10
  match afi ipv4
  match safi unicast
  match as-path "^20313_.*13$"
7100 Chassis(su-config-route-map-bgp)->
```

See [Chapter 39, Route-Map Manager Configuration](#) for BGP route-map configuration details.

## Route Selection Preference

When a route to the same prefix can be reached through two or more paths, the BGP local-preference attribute can be used within an AS to favor a specified path based on highest local-preference value. The local-preference value is only applicable within the local AS.

Use the **bgp local-pref** command, in BGP configuration mode, to set the local-preference for advertised routes.

The following example sets the local-preference for advertised routes to 150:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp router-id 151.1.1.9
7100 Chassis(su-config-bgp)->bgp local-pref 150
7100 Chassis(su-config-bgp)->
```

## Multi-Exit Discriminator (MED)

MEDs are used when an AS has multiple connections to another AS. Routes are advertised on both connections with different MEDs to specify a preferred path, typically for purposes of load balancing. By default, the MED for routes from different Autonomous Systems to the same destination are not compared. When two routes to the same destination are received from peers in different Autonomous Systems, the MED feature allows you to specify whether to compare those MEDs. When choosing between these routes, assuming that nothing else makes one preferable to the other (such as a configured policy), the values of the differing MEDs are used to choose the route to use. In this comparison, the route with the lowest MED is preferred.

By default, BGP sorts paths based on the neighbor AS and MED so that paths are sorted the same way every time. This results in a deterministic best-path selection. BGP can be configured to not take the neighbor AS into consideration when comparing the MED for each path. In this case, BGP compares MEDs when multiple routes with differing MEDs are received from peers in different autonomous systems.

The route MED value is set using a BGP route-map applied to the route.

Routes without a configured MED are treated as having the best possible MED.

Use the **bgp deterministic-med** command, in BGP configuration mode, to enable the deterministic processing of MEDs based upon the neighbor AS and MED. Deterministic MED is the default behavior.

The following example disables BGP deterministic MED for BGP router **65151**:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->no bgp deterministic-med
7100 Chassis(su-config-bgp)->
```

Use the **bgp always-compare-med** command, in BGP configuration mode, to compare MEDs when multiple routes with differing MEDs are received from peers in different autonomous systems.

The following example enables the comparison of MEDs from different ASs:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp always-compare-med
7100 Chassis(su-config-bgp)->
```

## Route Aggregation

Route aggregation allows you to aggregate one or more specific routes into a single route. Aggregate routes are only created if a more specific route of the aggregate route exists in the BGP routing table. Aggregate route configuration options provide for:

- Creating and advertising the aggregate route while at the same time suppressing the advertisement of all the more specific routes for this aggregate through route summarization.
- Retaining the advertisement of the AS-Path information for the specific routes of the aggregate. The default behavior for an aggregate route is to suppress the AS-Path information for the specific routes of the aggregate. It may be desirable to retain AS-Path information for routes in the aggregate that belong to an AS outside of the AS in which the aggregate is created.
- Enabling both route summarization and advertisement of the AS-Path information for the specific routes of the aggregate.
- Creating and advertising an aggregate while at the same time suppressing only those specific routes that match clauses in the applied route-map. Prefixes contained in the aggregate route that are not specifically matched in the route-map are not suppressed. You can not use this option in conjunction with route summarization.
- Creating and advertising an aggregate, while at the same time allow specifying in a route-map which AS path information is retained in the aggregate. This option is used in conjunction with retaining the advertisement of the AS-Path information for specific routes of the aggregate. You can not use this option in conjunction with route summarization.
- Creating and advertising an aggregate, while at the same time allowing for the modifying of aggregate route attributes specified in the route-map.

The following example creates and advertises aggregate route **200.51.0.0/22** and suppresses the advertisement of all the more specific routes for this aggregate:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->aggregate-address 200.51.0.0/22 summary
7100 Chassis(su-config-bgp)->
```

The following example sets the MED attribute to **50** for routes in aggregate route **200.51.0.0/22** using route-map **attrmap1**:

```
7100 Chassis(su-config)->route-map bgp attrmap1 permit 10
7100 Chassis(su-config-route-map-bgp)->set med 50
7100 Chassis(su-config-route-map-bgp)->exit
7100 Chassis(su-config)->show route-map attrmap1
    route-map bgp attrmap1 permit 10
        set med 50
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->aggregate-address 200.51.0.0/22 attribute-map
attrmap1
7100 Chassis(su-config-bgp)->
```

The following example retains AS-path information for routes **200.51.1.0/24** using route-map **advmap1** in aggregate route **200.51.0.0/22**:

```
7100 Chassis(su-config)->ip prefix-list advlist1 permit seq 1 200.51.1.0/24
7100 Chassis(su-config)->route-map bgp advmap1
7100 Chassis(su-config-route-map-bgp)->match prefix-list advlist1
7100 Chassis(su-config-route-map-bgp)->exit
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->aggregate-address 200.51.0.0/22 advertise-map
advmap1
7100 Chassis(su-config-bgp)->
```

## Source IP Address Update to the Peer

By default, BGP sets the source IP address of the BGP message to the outgoing interface. BGP checks the source IP address of the received message against the configured address of the BGP peer. If there is a mismatch, BGP discards the message. For EBGP neighbors that are directly connected and there is not an alternative path to the peer, the default outgoing interface as source address is sufficient. If the connection fails, no alternative route is available anyway.

In the case of an IBGP connected peer with an alternative route, should the connection on the outgoing interface fail, the route is withdrawn. By setting the source address to a virtual interface, such as a loopback interface, because the source address of the route is still available when the connection fails and an alternative route is available, BGP reconverges and installs the alternative route. If multiple paths exist between the BGP routers, using a loopback interface as the neighbor's source address can add stability to the network.

Use the **neighbor update-source** command in BGP router configuration mode to specify an alternative source address instead of the default directly connected interface address as the source address advertised for this IBGP router.

The following example causes the TCP session to peer **168.192.50.5** to be established over the loopback interface **4.3.2.1**:

```
7100 Chassis(su-config)->router bgp 65151
```

```
7100 Chassis(su-config-bgp)->bgp router-id 1.1.1.1
7100 Chassis(su-config-bgp)->neighbor 168.192.50.5 remote-as 5
7100 Chassis(su-config-bgp)->neighbor 168.192.50.5 update-source 4.3.2.1
```

See “[Configuring Source IP Address Update](#)” on page 35-35 for a remote peer source IP address update configuration example.

## Scalability and the Peer Full Mesh Requirement

BGP requires that all internal peers must establish a peer relationship with each other, which is called a full mesh. Full mesh networks scale very poorly. This can result in a large routing table and a management nightmare. BGP provides two techniques for reducing the full mesh: confederations and route reflectors.

### Confederations

The confederations extension to BGP, defined in RFC 3065, provides for the configuration of AS confederations. An AS confederation is a collection of routers, belonging to one or more autonomous systems, advertised as a single AS number to BGP speakers that are not members of the confederation.

Each AS confederation has a confederation ID. A router member of the confederation advertises itself to non-confederation peers using the AS confederation ID. A router member of the confederation advertises itself to other confederation peers using its AS number. A confederation ID can be any value from **1** to **65535**.

Each router member of a confederation must identify its confederation member peers.

Confederation information in the AS path sent to a neighbor peer is included by default. Inclusion of confederation information in the AS path sent to a neighbor peer can be disabled. It is possible to have AS path segments that do not adhere strictly to the confederation standard. Strict confederation path standards can be enforced. Strict confederation path enforcement is disabled by default.

See [Figure 35-1](#) on page 35-3 for a depiction of a BGP confederation topology.

Use the **bgp confederation-id** command, in BGP configuration mode, to specify the confederation this router belongs to.

The following example configures the BGP router to be a member of BGP confederation **100**:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp router-id 151.1.1.9
7100 Chassis(su-config-bgp)->bgp confederation-id 100
7100 Chassis(su-config-bgp)->
```

Use the **neighbor confed-member** command, in BGP configuration mode, to configure the neighbor as a member of the router’s confederation.

The following example configures neighbor **200.51.1.1** as a member of this router’s confederation:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp router-id 159.1.1.9
7100 Chassis(su-config-bgp)->bgp confederation-id 100
7100 Chassis(su-config-bgp)->neighbor 200.51.1.1 confed-member
7100 Chassis(su-config-bgp)->
```

Use the **neighbor aggregate-confed** command, in BGP configuration mode, to enable the inclusion of confederation information in the AS path sent to this router’s peers.



The following example disables the inclusion of confederation information in the AS paths sent to this router's peers:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp router-id 159.1.1.9
7100 Chassis(su-config-bgp)->no neighbor 200.51.1.1 aggregate-confed
7100 Chassis(su-config-bgp)->
```

Use the **bgp strict-confeds** command, in BGP configuration mode, to enable BGP to drop AS-Paths with non-standard confederation segments.

The following example enables the strict-confeds feature on this router:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp strict-confeds
7100 Chassis(su-config-bgp)->
```

See “[Configuring BGP Confederations](#)” on page 35-37 for a BGP confederation configuration example.

## Route Reflection

Route reflection enables you to configure a BGP speaker as a route reflector which passes internally learned routes to a cluster of linked IBGP neighbors. The route reflector configured router advertises the routes it has learned from each linked client to the other linked clients in the AS. In a route reflection topology, the route reflector is the hub, and each client only peers with the the hub. This eliminates the full mesh requirement. You can configure one or more routers in the AS to be reflectors. Some or all of the other routers for the AS are configured as clients. Route reflection is defined in RFC 4456.

Route reflection clients only peer with the route reflector. Route reflection configuration only occurs on the route reflector, identifying each route reflection client.

Multiple route reflectors can be configured in an AS. Multiple route reflectors can belong to a single route reflection cluster. A route reflection cluster is identified by a unique ID. If only a single route reflector is configured for an AS, the cluster ID defaults to the router ID of the route reflector.

See Autonomous System B of [Figure 35-1](#) on page 35-3 for a depiction of a route reflection topology.

Use the **neighbor route-reflector-client** command, in BGP configuration mode, to identify each client for the route reflection cluster.

The following example specifies that the neighbor **168.192.50.5** is a client of route reflector **1.1.1.1**:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp router-id 1.1.1.1
7100 Chassis(su-config-bgp)->neighbor 168.92.50.5 remote-as 5
7100 Chassis(su-config-bgp)->neighbor 168.92.50.5 route-reflector-client
```

Use the **bgp cluster-id** command, in BGP configuration mode, to specify a unique route reflection cluster ID the route reflector(s) belong to.

The following example configures a cluster ID of 1.2.3.4 for router 1.1.1.1:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp router-id 1.1.1.1
7100 Chassis(su-config-bgp)->bgp cluster-id 1.2.3.4
7100 Chassis(su-config-bgp)->
```

See “[Configuring Route Reflection](#)” on page 35-40 for a route reflection configuration example.

## Outbound Route Filtering (ORF)

The ORF feature allows a BGP speaker to notify a peer, using route-refresh messages, of the prefixes, communities, or extended-communities the router is interested in receiving updates for. Instead of using an inbound route-map to filter a large set of routes for installation into the local-RIB, this feature allows the router to ask for the particular set they are interested in. ORF can be configured to send route-refresh messages to the peer, receive them from the peer, or both.

The ORF capability is configured in BGP router configuration mode. The peer must also support the ORF capability. Configuring the ORF capability results in the advertisement of the ORF capability OPEN message. If the peer recognizes this capability, the peer advertises the capability in its OPEN message, otherwise, the peer sends a notification with an unrecognized capability error code. In this case, the local router will resend its OPEN message without the capability advertised.

Use the **bgp orf comm-filter** command to configure ORF for community filtering on both the local router and the peer router.

This example configures BGP to send the ORF capability for community filtering for IPv4 unicast to the peer:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp router-id 151.1.1.9
7100 Chassis(su-config-bgp)->bgp orf ipv4 unicast comm-filter send
7100 Chassis(su-config-bgp)->
```

Use the **bgp orf extcomm-filter** command to configure ORF for extended community filtering on both the local router and the peer router.

This example configures BGP to send ORF capability for extended community filtering for IPv4 unicast:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp router-id 151.1.1.9
7100 Chassis(su-config-bgp)->bgp orf ipv4 unicast extcomm-filter send
7100 Chassis(su-config-bgp)->
```

Use the **bgp orf prefix-filter** command to configure ORF for prefix filtering on this router.

This example configures BGP to send the ORF capability for prefix filtering to the BGP peer for IPv4 unicast:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp router-id 151.1.1.9
7100 Chassis(su-config-bgp)->bgp orf ipv4 unicast prefix-filter send
7100 Chassis(su-config-bgp)->
```

## Conditional Advertisement

The conditional advertisement feature allows a service provider to advertise certain routes to a preferred subnet under normal operational conditions, while maintaining the ability to move its traffic to an alternative subnet should its preferred routes fail. The conditional advertisement feature uses two route-maps to achieve this capability:

- The **non-exist-map** route-map which contains match prefix-list clauses for the preferred route(s) used under normal operational conditions

- The **advertise-map** route-map which contains match prefix-list clauses for the alternative route(s) used only if a preferred route is not available

Should any route specified in the non-exist route-map fail (no longer exist), BGP advertises all the routes specified in the advertise route-map, otherwise routes in the advertise route-map are not advertised to the peer. The conditional advertisement feature can be used to reduce traffic within an AS.

To configure the conditional advertisement feature:

- Create two prefix lists: one used to match prefixes for the advertise route-map; a second to match prefixes for the non-exist route-map
- Create a BGP route-map with a match clause for the advertise prefix-list
- Create a BGP route-map with a match clause for the non-exist prefix-list
- Apply the two route-maps to the **neighbor advertise-map** command

The following example:

- Configures an advertise map prefix-list named **adv-list1** and assigns it to BGP route-map **adv-map1**, specifying **prefix 1.0.0.0/8** as the advertised prefix
- Configures a non-exist map prefix-list named **non-exist-list1** and assigns it to BGP route-map **non-exist-map1**, specifying prefix **2.0.0.0/8** as the non-exist map prefix
- Configures a BGP advertise map for neighbor **192.168.12.112** that assigns **adv-map1** as the advertise map route-map and **non-exist-map1** as the non-exist map route-map

```
7100 Chassis(su-config)->ip prefix-list adv-list1 permit seq 1 1.0.0.0/8
7100 Chassis(su-config)->route-map bgp adv-map1
7100 Chassis(su-config-route-map-bgp)->match adv-list1 adv-map1
7100 Chassis(su-config-route-map-bgp)->exit
7100 Chassis(su-config)->ip prefix-list non-exist-list1 permit seq 1 2.0.0.0/8
7100 Chassis(su-config)->route-map bgp non-exist-map1
7100 Chassis(su-config-route-map-bgp)->match non-exist-list1 non-exist-map1
7100 Chassis(su-config-route-map-bgp)->exit
7100 Chassis(su-config)->router bgp 1
7100 Chassis(su-config-bgp)->neighbor 192.168.12.112 advertise-map adv-map1
non-exist-map non-exist-map1
7100 Chassis(su-config-bgp)->
```

See “[Configuring Conditional Advertisement](#)” on page 35-43 for a BGP conditional advertisement example.

## BGP Soft Reset

BGP soft reset configuration determines BGP behavior when inbound route policy changes. New policy needs to be applied to routes to determine which are admitted into the RIB. BGP soft reset can be applied in three ways.

### Internally Stored Route Reconfiguration

The soft reset of BGP routes can be based upon local internal storing of routes. For the soft reconfiguration approach, the received routes are stored in a dedicated table. When the inbound policy changes, the new policy is automatically applied. This approach is memory intensive, due to the storage of all the routes per peer or peer-group in a separate table.

Use the **neighbor soft-reconfiguration** command, in BGP router configuration mode, to enable the soft reconfiguration option on the local router. This command requires that you either specify the peer IP address or the peer group to be enabled for soft reconfiguration.

The following example turns on the route refresh capability for peer **10.10.25.1/24**:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp router-id 159.1.1.9
7100 Chassis(su-config-bgp)->neighbor 10.10.25.1/24 remote-as 5
7100 Chassis(su-config-bgp)->neighbor 10.10.25.1/24 soft-reconfiguration
```

## Route-Refresh

Route refresh, defined in RFC 2918, is a capability that the peer advertises in the OPEN message during session establishment. If both the local router and its peer agree to support this capability, a router can send a route refresh message to its peer whenever the local router's inbound policy changes. The peer will respond by resending its update messages. The local router can then reapply its policy without tearing down the BGP connection and without locally storing the received routes.

If the route refresh capability is enabled, the route refresh message is generated automatically when the inbound policy changes. Route refresh is enabled by default. If the soft reconfiguration soft reset method is enabled (see [Internally Stored Route Reconfiguration](#)) route refreshes are not sent.

Use the **bgp automatic-route-refresh** command, in BGP router configuration mode, to enable route refresh on the local router.

The following example disables BGP automatic route refresh for BGP router 65151:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->no bgp automatic-route-refresh
7100 Chassis(su-config-bgp)->
```

## Tear Down the BGP Connection

If neither route refresh nor soft reconfiguration are enabled, the only other method for dealing with soft refresh is to completely tear down the BGP connection. This approach is the most disruptive and requires reinitiation of the TCP and BGP connections and the reexchange of update messages. The reexchange of updates ensures that the new policy is applied to the routes. This approach is considered a last resort.

Use the **clear ip bgp** command, in any command mode, specifying the IP address of the peer to be torn down or an asterisk (\*) for all peers on this local router. Specifying the **soft** keyword sends a route refresh, if supported by the effected peers.

The following example clears the BGP peer 1.2.3.4:

```
7100 Chassis(rw)->clear ip bgp 1.2.3.4
```

The following example clears all BGP peers and sends a route refresh message to each cleared peer.

```
7100 Chassis(rw)->clear ip bgp * soft
```

## Community and Extended Community Attributes

BGP community and extended community attributes are optional, transitive BGP attributes that provide an administrative route labeling capability using route-maps. See [Chapter 39, Route-Map Manager Configuration](#) for details on configuring BGP route-maps.

## Community Attribute

Communities provide a label to a set of prefixes that share one or more common properties. Upstream providers use the community label to apply routing policy using route-maps to the community member prefixes.

Route-maps support the community attribute for both match and set clauses. Use the set community route-map clause to label the route as a member of the specified community.

There are two means of specifying a community name:

- Specify the AS number this route belongs to, followed by a colon (:), followed by the community number
- Specify a predefined community value, as defined in RFC 1997 and RFC 3765, that are supported by the community field such as:
  - **NO\_EXPORT** – Routes must not be advertised outside a BGP confederation boundary
  - **NO\_ADVERTISE** – Routes must not be advertised to other BGP peers
  - **NO\_EXPORT\_SUBCONFED** – Routes must not be advertised to external BGP peers
  - **NO\_PEER** – Routes must not be advertised across bilateral peer connections

One of the following actions must be specified:

- **remove-all** – Specifies that the action is to remove all communities from the route.
- **remove-specific** – Specifies that the action is to remove all matching communities from the route.
- **set-specific** – Specifies that the action is to append the specified community to the route.
- **remove-all-and-set** – Specifies that the action is to replace any existing communities in the route with the specified community

This example shows how to append the community value **100:100** to BGP routes matching prefix list named **permit100**:

```
7100 Chassis(su)->configure
7100 Chassis(su-config)->route-map bgp bgprml permit
7100 Chassis(su-config-route-map-bgp)->match prefix-list permit100
7100 Chassis(su-config-route-map-bgp)->set community 100:100 set-specific
7100 Chassis(su-config-route-map-bgp)->
```

This example shows how to append the well-known **NO\_PEER** community (RFC-3765) to BGP routes matching prefix list named **permit200**:

```
7100 Chassis(su)->configure
7100 Chassis(su-config)->route-map bgp bgprml permit
7100 Chassis(su-config-route-map-bgp)->match prefix-list permit200
7100 Chassis(su-config-route-map-bgp)->set community NO_PEER set-specific
7100 Chassis(su-config-route-map-bgp)->
```

The route-map **match community** clause provides the ability to set route policy for packets that have been set with community name matching the community name specified in the match clause.

This example shows how to match a packet community to community **100** in AS **121**:

```
7100 Chassis(su)->configure
7100 Chassis(su-config)->route-map bgp bgprml permit
7100 Chassis(su-config-route-map-bgp)->match community 121:100
```

```
7100 Chassis(su-config-route-map-bgp)->
```

## Extended Community Attribute

The Extended Community Attribute provides a mechanism for labeling information carried in BGP. It provides two important enhancements over the existing BGP Community Attribute:

- An extended range of use, ensuring that communities can be assigned to many different non-overlapping uses.
- The ability to specify a community type, providing structure for the community space. The 7100-Series supports extended community types:
  - IP, AS, and AS 4-octet route target
  - IP, AS, and AS 4-octet site-of-origin
  - OSPF domain ID, router ID, and route type

Use the set clause for the appropriate extended community type to label a route with the specified extended community. The extended community is labeled with a set value appropriate to the extended community. For example: The IPv4 route-target extended community requires a set value consisting of a valid IPv4 address followed by a colon (:) followed by a number in the range 0 - 65535. The firmware converts this set value into a hex identifier. The hex identifier for each set extended community is displayed in the **show ip bgp** output. When configuring an extended community match clause, use the **show ip bgp** command to determine the appropriate extended community identifier.

This example shows how to remove all matching extended communities from the AS route target **1001:10000** when all match clauses match for route-map **bgprm1**:

```
7100 Chassis(su)->configure
7100 Chassis(su-config)->route-map bgp bgprm1 permit
7100 Chassis(su-config-route-map-bgp)->match prefix-list permit200
7100 Chassis(su-config-route-map-bgp)->set extended-community as-route-target
1001:10000 remove-specific
7100 Chassis(su-config-route-map-bgp)->
```

This example shows how to match a packet against the extended community AS route target attribute **000203E9000186A0**:

```
7100 Chassis(su)->show ip bgp 1.0.0.0/8 detail
Route status codes: > - active
```

Network	Next Hop	Rib	MED	Local-Pref	Origin	AS	Path
> 1.0.0.0/8	192.168.121.112	U	0	100	IGP	121	2013

```
Community attributes in route:
```

```
121:100
```

```
Extended Community attributes in route:
```

```
Route Target: 1001:10000 (0x000203E9000186A0)
```

```
7100 Chassis(su)->configure
7100 Chassis(su-config)->route-map bgp bgprm1 permit
7100 Chassis(su-config-route-map-bgp)->match extended-community 000203E9000186A0
7100 Chassis(su-config-route-map-bgp)->
```

## Route Flap Dampening

BGP route flap dampening is used to suppress routes that have been unstable due to misconfiguration, a rebooting router or module, or link flapping. For each route flap, a penalty is assessed to the route. When this route penalty reaches the cutoff threshold, the route is suppressed. The route penalty decays over time based upon a configured decay half-life value. Once the decaying penalty reaches the reuse threshold, or the hold-time timer has expired, the route is reinstalled into the routing table.


The route penalty starts with a value of **0**. With each route instability, the route penalty is increased by **100**. Route flap dampening provides for the following thresholds and timers that interact with the route penalty:

- **Cutoff threshold** – Specifies the route penalty value beyond which the route is suppressed.
- **Half-life reachable timer** – Specifies the time, in seconds, it takes a route penalty to decay to half of its current value, assuming the route is both reachable and remains stable during that period
- **Half-life unreachable timer** – Specifies the time, in seconds, it takes a route penalty to decay to half of its current value, assuming the route is both unreachable and remains stable during that period
- **Memory limit reachable timer** - Specifies the maximum time, in seconds, the history of a previous instability is retained in memory for a reachable route
- **Memory limit unreachable timer** - Specifies the maximum time, in seconds, the history of a previous instability is retained in memory for an unreachable route
- **Reuse threshold** – Specifies the route penalty below which a suppressed route is reused (unsuppressed).
- **Hold-time timer** – Specifies the maximum time a route can be suppressed regardless of its stability history

Route flap dampening configuration is contained in an administratively named flap table initially configured with default timer and threshold values. The flap table is assigned to a route-map. The route-map is assigned to the neighbor address to be monitored as an inbound route-map. When the first route instability occurs, the route penalty for this flap table is set to 100. This penalty immediately starts to decay at the rate set by the appropriate half-life timer. Upon the occurrence of a second route instability, the current route penalty is increased by 100. If the current penalty is now greater than the cutoff threshold (default penalty of 125) the route is suppressed. If no further instability occurs, the route penalty will eventually decay below the value of the reuse threshold, based upon the appropriate half-life timer setting. When the route penalty falls below the value of the reuse threshold, the route is unsuppressed. If route instability continues to occur, adding penalty points to the route penalty at a greater rate than the penalty can decay, it is possible for a route to stay suppressed until the hold-time timer expires. Once the hold-time timer expires, the route is unsuppressed, regardless of the current route penalty.

[Figure 35-2](#) presents a default route flap dampening timing example. The route penalty starts at 0 and stays there until the first route flap. At the first route flap, route flap dampening sets the route penalty to 100. It immediately starts to decay based upon the default reachable half-life of 300 seconds. Approximately 20.5 seconds later, a second flap occurs. Route flap dampening adds 100 to the current route penalty of 96. The current route penalty is now greater than the cutoff value. The route is suppressed. No further route flaps occur.

### Figure 35-2 Route Flap Dampening Timing



The route penalty immediately starts to decay. 300 seconds later the first reachable half-life is reached and the penalty is now 98. Because the route penalty is still greater than the reuse setting, route flap dampening continues to suppress the route. Route flap dampening unsuppresses the route when the route penalty decays to 50, the default reuse setting. Because no further route flaps occur, the route penalty continues to decay until it reaches 0.

The memory limit timers are used by route flap dampening for internal calculations. Half-life timers must be configured to a value less than the corresponding reachable or unreachable memory limit timer.

The flap table flap count or all flap statistics can be cleared. When clearing the flap count on a suppressed route, the route remains suppressed. When clearing all statistics on a suppressed route, the route is unsuppressed, regardless of the current route penalty value.

Use the **dampen-flap** command, in router configuration mode, to name the flap table and enter route flap dampening configuration mode.

The following example enters route flap dampening configuration mode for flap table **flap1**.

```
7100 Chassis(su-config)-> dampen-flap flap1
7100 Chassis(su-config-dampen-flap)->
```

Use the **cutoff** command, in route flap dampening configuration mode, to modify the route suppression threshold for this flap table context.

The following example modifies the cutoff threshold from the default value of 125 to **150** for the flap table **flap1**.

```
7100 Chassis(su-config)->dampen-flap flap1
7100 Chassis(su-config-dampen-flap)->cutoff 150
7100 Chassis(su-config-dampen-flap)->
```



Use the **half-life-reach** command, in route flap dampening configuration mode, to specify the time in seconds after which a reachable route's penalty value decays to half of its current value, assuming no further route instability for this route.

The following example configures the half life reachable value to **250** seconds for the **flap1** flap table.

```
7100 Chassis(su-config)->dampen-flap flap1
7100 Chassis(su-config-dampen-flap)->half-life-reach 250
7100 Chassis(su-config-dampen-flap)->
```

Use the **half-life-unreach** command, in route flap dampening configuration mode, to specify the time in seconds after which an unreachable route's penalty value decays to half of its current value, assuming no further route instability for this route.

The following example configures the half life unreachable value to **600** seconds for the **flap1** flap table.

```
7100 Chassis(su-config)->dampen-flap flap1
7100 Chassis(su-config-dampen-flap)->half-life-unreach 600
7100 Chassis(su-config-dampen-flap)->
```

Use the **hold-time** command, in route flap dampening configuration mode, to specify the maximum amount of time, in seconds, a route can be suppressed for this route table context.

The following example configures the hold-time to **1000** seconds for flap table **flap1**:

```
7100 Chassis(su-config)->dampen-flap flap1
7100 Chassis(su-config-dampen-flap)->hold-time 1000
```

Use the **reuse** command, in route flap dampening configuration mode, to specify the route penalty threshold under which a suppressed route is reused (unsuppressed).

The following example configures the reuse threshold to the route penalty value of **75** for flap table **flap1**:

```
7100 Chassis(su-config)->dampen-flap flap1
7100 Chassis(su-config-dampen-flap)->reuse 75
```

Use the **memory-limit-reach** command, in route flap dampening configuration mode, to specify the maximum time, in seconds, any memory of a previous instability is retained for this route table.

The following example configures the memory limit reachable value to **800** seconds for the **flap1** flap table.

```
7100 Chassis(su-config)->dampen-flap flap1
7100 Chassis(su-config-dampen-flap)->memory-limit-reach 800
7100 Chassis(su-config-dampen-flap)->
```

Use the **memory-limit-unreach** command, in route flap dampening configuration mode, to specify the maximum time, in seconds, any memory of a previous instability is retained for this route table.

The following example configures the memory limit unreachable value to **1600** seconds for the **flap1** flap table.

```
7100 Chassis(su-config)->dampen-flap flap1
7100 Chassis(su-config-dampen-flap)->memory-limit-unreach 1600
7100 Chassis(su-config-dampen-flap)->
```

See "[Configuring Flap Dampening](#)" on page 35-46 for a BGP flap dampening configuration example.

## Graceful Restart

BGP graceful restart provides for the continued processing and packet forwarding of a router's data-forwarding plane even if the router control plane fails. With both a router and its peer graceful restart enabled, BGP exchanges the graceful restart capability (BGP code 64) in the initial BGP OPEN messages that establish the session.

When a failure takes place and the router restarts its BGP process, normally peer routers clear all routes associated with the restarting router. When graceful restart is enabled on a router, the peer router marks all routes as "stale" and continues to forward packets based on the expectation that the restarting router will reestablish the BGP session within a reasonable period of time. During the period of the restart, the restarting router continues to forward packets based upon routing state at the time of the restart. Peers refresh the restarting router with RIB updates.

When the restarting router opens the new BGP session, it will again send the BGP capability code 64 to its peers with flags set to let the peer router know that the BGP process has restarted. When the restarting router completes its restart and RIB update, it in turn updates its peers with any new updates.

Graceful restart reduces routing flaps, which stabilizes the network and reduces the consumption of control-plane resources.

BGP graceful restart timing is based upon four configurable intervals:

- **Restart defer interval** – Specifies the upper bound (in seconds) on the amount of time route selection will be deferred when BGP is restarting. The value specified should be large enough to provide all peers with enough time to send all their routes. The value must be greater than or equal to the restart timeout setting.
- **Restart timeout interval** – Specifies the interval which BGP advertises to its peers, in the OPEN message exchange, as the estimated time (in seconds) it will take for the BGP session to be reestablished after a restart. This can be used to speed up routing convergence by its peer in case the BGP speaker does not come back after a restart. Following a local restart, BGP will impose the restart timeout value as the upper bound on the length of time permitted for BGP to restart. If BGP fails to restart within the restart timeout period, route selection will commence immediately thereby overriding the restart defer timer.
- **Restart time interval** – Allows the peer to configure the maximum time (in seconds) it will wait for the restarting router to come back after a restart. This value will be used instead of the restart timeout value advertised in the OPEN message exchange, if the OPEN message value exceeds this restart timer value.
- **Stale path interval** – Configures the maximum time following a restart before removing stale routes from the peer. The stale path interval must be greater than or equal to the restart time.

Graceful restart must be enabled for the four configurable graceful restart timers to be relevant.

Use the **bgp graceful-restart** command, in BGP configuration mode, to enable graceful restart on the local router.

The following example enables graceful restart on router 151.1.1.9

```
7100 Chassis(su-config-bgp)->bgp router 65151
7100 Chassis(su-config-bgp)->bgp router-id 151.1.1.9
7100 Chassis(su-config-bgp)->bgp graceful-restart
7100 Chassis(su-config-bgp)->
```

Use the **bgp restart-defer** command, in BGP configuration mode, to configure the time to defer route selection after gracefully restarting.

The following example configures the defer timer to 150 seconds:

```
7100 Chassis(su-config)->router bgp 65151
```

```
7100 Chassis(su-config-bgp)->bgp router-id 151.1.1.9
7100 Chassis(su-config-bgp)->bgp restart-defer 150
7100 Chassis(su-config-bgp)->
```

Use the **bgp restart-time** command, in BGP configuration mode, to configure the maximum time to wait for a graceful restarting peer to come back up after a restart.

The following example configures the restart time to be 100 seconds:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp router-id 159.1.1.9
7100 Chassis(su-config-bgp)->bgp restart-time 100
7100 Chassis(su-config-bgp)->
```

Use the **bgp restart-timeout** command, in BGP configuration mode, to configure the estimated time advertised to peers in the OPEN message for the session to be reestablished after a graceful restart.

The following example configures the restart-timeout to be 150 seconds:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp router-id 159.1.1.9
7100 Chassis(su-config-bgp)->bgp restart-timeout 150
```

Use the **bgp stale-path-time** command, in BGP configuration mode, to configure the maximum time following a restart stale routes are allowed to persist on the peer.

The following example sets the stale-path-time to 150 seconds:

```
7100 Chassis(su-config)->router bgp 65151
7100 Chassis(su-config-bgp)->bgp router-id 159.1.1.9
7100 Chassis(su-config-bgp)->bgp stale-path-time 150
```

## Configuring BGP

For information about...	Refer to page...
<a href="#">Configuring Basic BGP Router Parameters</a>	35-25
<a href="#">Configuring BGP Route Injection</a>	35-26
<a href="#">Configuring External BGP Basic Peering</a>	35-27
<a href="#">Configuring Internal BGP Basic Peering</a>	35-29
<a href="#">Configuring Multihop EBGp Basic Peering</a>	35-31
<a href="#">Configuring BGP Neighbor Parameters</a>	35-34
<a href="#">Configuring Source IP Address Update</a>	35-35
<a href="#">Configuring BGP Confederations</a>	35-37
<a href="#">Configuring Route Reflection</a>	35-40
<a href="#">Configuring Outbound Route Filtering (ORF)</a>	35-43
<a href="#">Configuring Conditional Advertisement</a>	35-43
<a href="#">Configuring BGP Soft Reset</a>	35-46
<a href="#">Configuring Flap Dampening</a>	35-46

For information about...	Refer to page...
<a href="#">Configuring Graceful Restart</a>	35-50
<a href="#">BGP Monitoring and Clearing</a>	35-51

This section provides details for BGP configuration on 7100-Series products.

[Table 35-2](#) lists BGP default values.

**Table 35-2 Default BGP Parameters**

Parameter	Description	Default Value
advertisement interval	The minimum interval in seconds between sending EBGP routing updates.	30 seconds
AS origination interval	The interval in seconds between successive update messages for route prefixes that originate in the local AS.	15 seconds
AS path limit	The upper limit on the AS path length when configuring a route.	1
connection retry interval	The amount of time between attempts to reestablish a connection to configured peers that are no longer available.	120 seconds
cutoff	The route suppression threshold used by flap dampening to determine when a flapping route should be suppressed.	125
distance (External to AS)	The priority given to external BGP routes relative to other protocols for the local router.	20
distance (Internal to AS)	The priority given to internal BGP routes relative to other protocols for the local router.	200
graceful restart	A BGP extension that provides for the continued processing and forwarding of packets by the data-forwarding plane even if the control plane fails.	disabled
half-life reachable	The time in seconds after which a reachable route's penalty value decays to half of its current value.	300 seconds
half-life unreachable	The time in seconds after which an unreachable route's penalty value decays to half of its current value.	900 seconds
hold-time (flap dampening)	The maximum amount of time a route can be suppressed.	900 seconds
hold-time (peering session negotiation)	The number of seconds to use when negotiating a peering session within a group.	90 seconds

**Table 35-2 Default BGP Parameters (continued)**

Parameter	Description	Default Value
idle hold interval	The interval in seconds between returning to the idle state and reinitiating a TCP connection for the peer.	15 seconds
keepalive timer	The interval between keepalive messages	30 seconds or one-third of the hold-time setting
local preference	The preference for this route over other possible routes on the local router.	100
maximum EBGp ECMP routes	The maximum number of external BGP ECMP routes on the local router.	1
maximum IBGP ECMP routes	The maximum number of internal BGP ECMP routes on the local router.	1
maximum ORF entries	The maximum number of outbound route filtering entries that will be accepted from the peer.	100000
maximum prefixes	The peak number of prefixes that BGP will accept for installation into the routing information base.	0 – unlimited
MED	The Multi-Exit Discriminator value when configuring a route.	0
memory limit reachable	The maximum time in seconds any memory of a previous instability is retained for a reachable route, given the route state is both unchanged and reachable.	700 seconds
memory limit unreachable	The maximum time in seconds any memory of a previous instability is retained for an unreachable route, given the route state is both unchanged and unreachable.	1800 seconds
open delay	The interval in seconds between the establishment of a TCP connection and the sending of an OPEN message to open a BGP session.	0 – no delay
origin	The value of the origin process attribute when configuring a route	0 – IGP
peer type	The type of peer or peer group	IBGP
peering type	Determines whether updates for prefixes containing the NOPEER community will be accepted by or sent to this neighbor.	unspecified

**Table 35-2 Default BGP Parameters (continued)**

Parameter	Description	Default Value
restart defer period	The time in seconds that route selection is deferred after a graceful restart.	120 seconds
restart time	The time in seconds to wait for a graceful restart capable peer to come back after a graceful restart.	120 seconds
restart timeout	The estimated time in seconds that is advertised to peers in the OPEN message for the session to be reestablished after a graceful restart.	120
reuse penalty	The route penalty value below which a suppressed route is reused.	50
route withdrawal interval	The interval between the advertisement and subsequent withdrawal of a route.	30 seconds
stale path time	The maximum time in seconds following a restart before removing stale routes from the peer.	120
time-to-live (TTL)	Specifies the number of hops for this neighbors TTL.	64

## Configuring Basic BGP Router Parameters

The basic steps for configuring BGP are:

- Entering BGP configuration mode for this router, specifying the AS
- Setting a BGP-specific router ID
- Configuring BGP parameters

[Procedure 35-1](#) describes how to configure Basic BGP.

### Procedure 35-1 Configuring Basic BGP

Step	Task	Command(s)
1.	In configuration command mode, enable BGP and enter BGP configuration mode, specifying the autonomous system for this router.	<b>router bgp</b> <i>as-number</i>
2.	In BGP configuration mode, configure the BGP router-ID.	<b>bgp router-id</b> <i>router-id</i>
3.	Optionally enter address family mode and configure the address family indicator (AFI) for BGP peers.	<b>address-family</b> [ <i>ipv4</i>   <i>ipv6</i> ] [ <b>unicast</b>   <b>multicast</b>   <b>both</b> ] [ <i>vrf vrf-name</i> ]
4.	In BGP configuration mode, optionally configure an aggregate by combining the characteristics of multiple routes so that a single route is advertised.	<b>aggregate-address</b> <i>prefix/length</i> [ <b>summary</b> ] [ <b>as-set</b> ] [ <b>summary-and-as-set</b> ] [ <b>suppress-map</b> <i>route-map</i> ] [ <b>advertise-map</b> <i>route-map</i> ] [ <b>attribute-map</b> <i>route-map</i> ]

**Procedure 35-1 Configuring Basic BGP (continued)**

Step	Task	Command(s)
5.	In BGP configuration mode, optionally enable aggregation of routes independent of the route MED.	<b>bgp aggregate-med</b>
6.	In BGP configuration mode, optionally disable deterministic processing of MEDs.	<b>no bgp deterministic-med</b>
7.	In BGP configuration mode, optionally specify whether to compare MEDs when multiple routes with differing MEDs are received from peers in different Autonomous Systems.	<b>bgp always-compare-med</b>
8.	In BGP configuration mode, optionally modify the local-preference of advertised routes for the router.	<b>bgp local-pref</b> <i>pref-value</i>
9.	In BGP configuration mode, optionally modify the route selection priority given to internal or external BGP routes compared to other protocols for the router.	<b>bgp distance</b> { <b>internal</b>   <b>external</b> } <i>distance</i>
10.	In BGP configuration mode, optionally modify the maximum number of allowed external BGP ECMP routes.	<b>bgp max-ebgp-ecmp-routes</b> <i>value</i>
11.	In BGP configuration mode, optionally modify the maximum number of allowed internal BGP ECMP routes.	<b>bgp max-ibgp-ecmp-routes</b> <i>value</i>
12.	In BGP configuration mode, optionally disable BGP configuration on the router. Administratively disabled BGP configuration can be reenabled using the <b>enable</b> command.	<b>no enable</b>
13.	In BGP configuration mode, optionally disable message logging via the syslog mechanism whenever a BGP peer enters or leaves the established state.	<b>no log-up-down</b>
14.	In BGP configuration mode, optionally enable the sending of BGP traps when a peer transitions to Established or a lower state.	<b>bgp trap</b> { <b>peer-established</b>   <b>peer-degraded</b> }

## Configuring BGP Route Injection

Routes can be injected into BGP by route redistribution or by specifying the network prefixes to import into BGP.

[Procedure 35-2](#) describes how to inject routes into BGP.

**Procedure 35-2 Configuring BGP Route Injection**


Step	Task	Command(s)
1.	In BGP configuration mode, optionally specify network prefixes to be imported into BGP.	<b>network</b> <i>prefix/length</i> [ <b>route-map</b> <i>name</i> ][ <b>aspath-limit</b> <i>limit</i> ] [ <b>origin code</b> ] [ <b>med value</b> ] [ <b>local-pref value</b> ]

**Procedure 35-2 Configuring BGP Route Injection (continued)**

Step	Task	Command(s)
2.	In BGP configuration mode, optionally specify that connected routes are redistributed into BGP.	<b>redistribute connected</b> [aspath-limit <i>limit</i> ] [origin <i>code</i> ] [med <i>value</i> ] [local-pref <i>value</i> ] [route-map <i>name</i> ]
3.	In BGP configuration mode, optionally specify that RIP routes are redistributed into BGP.	<b>redistribute rip</b> [aspath-limit <i>limit</i> ] [origin <i>code</i> ] [med <i>value</i> ] [local-pref <i>value</i> ] [route-map <i>name</i> ]
4.	In BGP configuration mode, optionally specify that static routes are redistributed into BGP.	<b>redistribute static</b> [aspath-limit <i>limit</i> ] [origin <i>code</i> ] [med <i>value</i> ] [local-pref <i>value</i> ] [route-map <i>name</i> ]
5.	In BGP configuration mode, optionally specify that OSPF routes are redistributed into BGP.	<b>redistribute ospf</b> <i>proc-id</i> [aspath-limit <i>limit</i> ] [origin <i>code</i> ] [med <i>value</i> ] [local-pref <i>value</i> ] [route-map <i>name</i> ]

**Configuring External BGP Basic Peering**

The following example configures a basic External BGP peering as displayed in [Figure 35-3](#) on page 35-27.

**Figure 35-3 Basic EBGP Peering Topology**

The example consists of two routers with a connection on subnet 192.168.12.0/24. One router belongs to AS 1. The other router belongs to AS 2. The example configuration consists of configuring, for each router:

- The connection interface and IP address
- The AS the router belongs to
- The router ID
- The connection neighbor IP address and remote AS
- The redistribution of static routes

**Router 1**

```
Router 1 (rw) -> configure
Router 1 (rw-config) -> interface vlan 12
```



```
Router 1 (rw-config-intf-vlan.0.12) -> ip address 192.168.12.111 255.255.255.0
Router 1 (rw-config-intf-vlan.0.12) -> no shutdown
Router 1 (rw-config-intf-vlan.0.12) -> exit
```

```
Router 1 (rw) -> configure
Router 1 (rw-config) -> router bgp 1
Router 1 (su-config-bgp) -> bgp router-id 1.1.1.1
Router 1 (su-config-bgp) -> neighbor 192.168.12.112 remote-as 2
Router 1 (su-config-bgp) -> redistribute static
Router 1 (su-config-bgp) ->
```

## Router 2

```
Router 2 (rw) -> configure
Router 2 (rw-config) -> interface vlan 12
Router 2 (rw-config-intf-vlan.0.12) -> ip address 192.168.12.112 255.255.255.0
Router 2 (rw-config-intf-vlan.0.12) -> no shutdown
Router 2 (rw-config-intf-vlan.0.12) -> exit
```

```
Router 2 (rw) -> configure
Router 2 (rw-config) -> interface vlan 13
Router 2 (rw-config-intf-vlan.0.13) -> ip address 192.168.13.111 255.255.255.0
Router 2 (rw-config-intf-vlan.0.13) -> no shutdown
Router 2 (rw-config-intf-vlan.0.13) -> exit
```

```
Router 2 (rw) -> configure
Router 2 (rw-config) -> router bgp 2
Router 2 (su-config-bgp) -> bgp router-id 2.2.2.2
Router 2 (su-config-bgp) -> neighbor 192.168.12.111 remote-as 1
Router 2 (su-config-bgp) -> redistribute static
Router 2 (su-config-bgp) ->
```

[Procedure 35-3](#) is a simple configuration intended for external BGP propagation.

### Procedure 35-3 EBGW Basic Peering Configuration

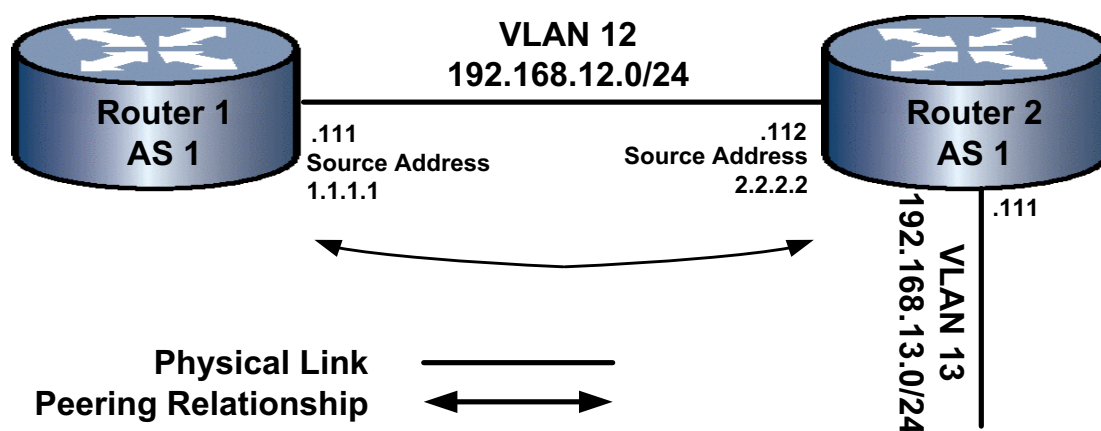
Step	Task	Command
1.	In configuration mode, configure static routes between BGP routers to allow IP traffic transmission between remote routers.	<b>ip route</b> { <i>prefix mask</i>   <i>prefix/prefix-length</i> } { <i>ip-address</i> [ <b>recursive</b>   <b>interface</b> <i>interface-name</i> ]   <b>interface</b> <i>interface-name</i>   <b>vlan</b> <i>vlan-id</i>   <b>vrf</b> <i>egress-vrf</i>   <b>blackhole</b>   <b>reject</b> } [ <i>distance</i> ] [ <b>tag</b> <i>tag-id</i> ]
2.	In configuration mode, configure loopback and physical addresses and enter interface configuration mode.	<b>interface</b> { <b>vlan</b> <i>vlan-id</i>   <b>loopback</b> <i>loopback-id</i>   <i>interface-name</i> }
3.	In interface configuration mode, configure the IP address for the interface that serves as the BGP speaker.	<b>ip address</b> { <i>ip-address ip-mask</i>   <i>ip-address/prefixLength</i> } [ <b>primary</b>   <b>secondary</b> ]

**Procedure 35-3 EBGP Basic Peering Configuration (continued)**

Step	Task	Command
4.	In configuration mode, specify an AS number for the router and enter BGP Configuration mode.	<b>router bgp</b> <i>as-number</i>
5.	In BGP configuration mode, configure a BGP-specific router ID to override the global router ID.	<b>bgp router-id</b> <i>router-id</i>
6.	In BGP configuration mode, configure the peer by identifying its IP address and AS.	<b>neighbor</b> <i>ip-address</i> <b>remote-as</b> <i>as-num</i> [ <b>password</b> <i>password</i> ]
7.	In BGP configuration mode, redistribute routes into BGP, optionally specifying a route-map. Supported <i>route-types</i> are connected, static, OSPF, and RIP.	<b>redistribute</b> <i>route-type</i> [ <b>aspath-limit</b> <i>limit</i> ] [ <b>origin</b> <i>code</i> ] [ <b>med</b> <i>value</i> ] [ <b>local-pref</b> <i>value</i> ] [ <b>route-map</b> <i>name</i> ]

## Configuring Internal BGP Basic Peering

The following example configures a basic internal BGP peering as displayed in [Figure 35-4](#) on page 35-29. The example configures the source IP address using the update source address feature to take advantage of the full mesh within an IBGP topology. See “[Source IP Address Update to the Peer](#)” on page 35-10 for an explanation of this feature.

**Figure 35-4 Basic IBGP Peering Topology**

The example consists of two routers with a connection on subnet 192.168.12.0/24. Because this is an internal BGP connection, both routers belong to the same AS. The example configuration consists of configuring, for each router:

- The loopback interface for the update source IP address
- The connection interface and IP address
- The AS the router belongs to
- The router ID
- The connection neighbor IP address (use update source IP address) and remote AS
- The connection neighbor IP address, specifying the update source IP address for this router
- The redistribution of static routes

## Router 1

```
Router 1 (rw)->configure
Router 1 (rw-config)->interface loopback 1
Router 1 (rw-config-intf-loop.0.1)->ip address 1.1.1.1 255.255.255.0
Router 1 (rw-config-intf-loop.0.1)->no shutdown
Router 1 (rw-config-intf-loop.0.1)->exit

Router 1 (rw-config)->interface vlan 12
Router 1 (rw-config-intf-vlan.0.12)->ip address 192.168.12.111 255.255.255.0
Router 1 (rw-config-intf-vlan.0.12)->no shutdown
Router 1 (rw-config-intf-vlan.0.12)->exit

Router 1 (rw)->configure
Router 1 (rw-config)->router bgp 1
Router 1 (su-config-bgp)->bgp router-id 1.1.1.1
Router 1 (su-config-bgp)->neighbor 2.2.2.2 remote-as 1
Router 1 (su-config-bgp)->neighbor 2.2.2.2 update-source 1.1.1.1
Router 1 (su-config-bgp)->redistribute static
Router 1 (su-config-bgp)->
```

## Router 2

```
Router 1 (rw)->configure
Router 1 (rw-config)->interface loopback 1
Router 1 (rw-config-intf-loop.0.1)->ip address 2.2.2.2 255.255.255.0
Router 1 (rw-config-intf-loop.0.1)->no shutdown
Router 1 (rw-config-intf-loop.0.1)->exit

Router 2 (rw-config)->interface vlan 12
Router 2 (rw-config-intf-vlan.0.12)->ip address 192.168.12.112 255.255.255.0
Router 2 (rw-config-intf-vlan.0.12)->no shutdown
Router 2 (rw-config-intf-vlan.0.12)->exit

Router 2 (rw)->configure
Router 2 (rw-config)->interface vlan 13
Router 2 (rw-config-intf-vlan.0.13)->ip address 192.168.13.111 255.255.255.0
Router 2 (rw-config-intf-vlan.0.13)->no shutdown
Router 2 (rw-config-intf-vlan.0.13)->exit

Router 2 (rw)->configure
Router 2 (rw-config)->router bgp 2
Router 2 (su-config-bgp)->bgp router-id 2.2.2.2
Router 2 (su-config-bgp)->neighbor 1.1.1.1 remote-as 1
Router 1 (su-config-bgp)->neighbor 1.1.1.1 update-source 2.2.2.2
Router 2 (su-config-bgp)->redistribute static
Router 2 (su-config-bgp)->
```

[Procedure 35-4](#) on page 35-31 is a simple configuration intended for internal BGP propagation.

#### Procedure 35-4 IBGP Basic Peering Configuration

Step	Task	Command
1.	In configuration mode, configure static routes between BGP routers to allow IP traffic transmission between remote routers.	<b>ip route</b> { <i>prefix mask</i>   <i>prefix/prefix-length</i> } { <i>ip-address</i> [ <b>recursive</b>   <b>interface</b> <i>interface-name</i> ]   <b>interface</b> <i>interface-name</i>   <b>vlan</b> <i>vlan-id</i>   <b>vrf</b> <i>egress-vrf</i>   <b>blackhole</b>   <b>reject</b> } [ <i>distance</i> ] [ <b>tag</b> <i>tag-id</i> ]
2.	In configuration mode, configure loopback and physical addresses and enter interface configuration mode.	<b>interface</b> { <b>vlan</b> <i>vlan-id</i>   <b>loopback</b> <i>loopback-id</i>   <i>interface-name</i> }
3.	In interface configuration mode, configure the IP address for the interface that serves as the BGP speaker.	<b>ip address</b> { <i>ip-address ip-mask</i>   <i>ip-address/prefixLength</i> } [ <b>primary</b>   <b>secondary</b> ]
4.	In configuration mode, specify an AS number for the router and enter BGP Configuration mode.	<b>router bgp</b> <i>as-number</i>
5.	In BGP configuration mode, configure a BGP-specific router ID to override the global router ID.	<b>bgp router-id</b> <i>router-id</i>
6.	In BGP configuration mode, specify an update source IP address assigned to a loopback interface, for this router, to be used as the source address instead of the default outgoing interface IP address.	<b>neighbor</b> <i>ip-address</i> <b>update-source</b> <i>source-addr</i>
7.	In BGP configuration mode, configure the peer by identifying its source IP address and AS.	<b>neighbor</b> <i>ip-address</i> <b>remote-as</b> <i>as-num</i>
8.	In BGP configuration mode, redistribute routes into BGP, optionally specifying a route-map. Supported <i>route-types</i> are connected, static, OSPF, and RIP.	<b>redistribute</b> <i>route-type</i> [ <b>aspath-limit</b> <i>limit</i> ] [ <b>origin</b> <i>code</i> ] [ <b>med</b> <i>value</i> ] [ <b>local-pref</b> <i>value</i> ] [ <b>route-map</b> <i>name</i> ]


## Configuring Multihop EBGW Basic Peering

An EBGW Multihop configuration is a topology where external BGP neighbors are not connected to the same subnet. Such neighbors are logically, but not physically connected. For example, BGP can be run between external neighbors across non-BGP routers.

Be aware that no IP traffic can pass to advertised BGP routes until an IGP protocol or static route is configured for those prefixes on the middle router.

See [Figure 35-5](#) on page 35-31 for a presentation of the multihop EBGW basic peering configuration example topology.

#### Figure 35-5 EBGW Multihop Peering Topology



### Router 1

```
Router 1 (rw) -> configure
Router 1 (rw-config) -> interface vlan 12
Router 1 (rw-config-intf-vlan.0.12) -> ip address 192.168.12.111 255.255.255.0
Router 1 (rw-config-intf-vlan.0.12) -> no shutdown
Router 1 (rw-config-intf-vlan.0.12) -> exit
```

```
Router 1 (rw) -> configure
Router 1 (rw-config) -> router bgp 1
Router 1 (su-config-bgp) -> bgp router-id 1.1.1.1
Router 1 (su-config-bgp) -> neighbor 192.168.13.111 remote-as 2
Router 1 (su-config-bgp) -> redistribute static
Router 1 (su-config-bgp) ->
```

### Router 2

```
Router 2 (rw) -> configure
Router 2 (rw-config) -> interface vlan 12
Router 2 (rw-config-intf-vlan.0.12) -> ip address 192.168.12.112 255.255.255.0
Router 2 (rw-config-intf-vlan.0.12) -> no shutdown
Router 2 (rw-config-intf-vlan.0.12) -> exit
```

```
Router 2 (rw-config) -> interface vlan 13
```

```
Router 2 (rw-config-intf-vlan.0.13) -> ip address 192.168.13.111 255.255.255.0
Router 2 (rw-config-intf-vlan.0.13) -> no shutdown
Router 2 (rw-config-intf-vlan.0.13) -> exit
```

### Router 3

```
Router 3 (rw) -> configure
Router 1 (rw-config) -> interface vlan 13
Router 1 (rw-config-intf-vlan.0.13) -> ip address 192.168.13.112 255.255.255.0
Router 1 (rw-config-intf-vlan.0.13) -> no shutdown
Router 1 (rw-config-intf-vlan.0.13) -> exit
```

```
Router 1 (rw) -> configure
Router 1 (rw-config) -> ip route 192.168.12.0/24 192.168.12.111 interface vlan 13
Router 1 (rw-config) -> router bgp 2
Router 1 (su-config-bgp) -> bgp router-id 2.2.2.2
Router 1 (su-config-bgp) -> neighbor 192.168.12.111 remote-as 1
Router 1 (su-config-bgp) -> redistribute static
Router 1 (su-config-bgp) ->
```

[Procedure 35-5](#) is a simple configuration intended for multihop BGP propagation.

### Procedure 35-5 Multihop BGP Basic Peering Configuration

Step	Task	Command
1.	In configuration mode, configure static routes between BGP routers to allow IP traffic transmission between remote routers.	<b>ip route</b> { <i>prefix mask</i>   <i>prefix/prefix-length</i> } { <i>ip-address</i> [ <b>recursive</b>   <b>interface</b> <i>interface-name</i> ]   <b>interface</b> <i>interface-name</i>   <b>vlan</b> <i>vlan-id</i>   <b>vrf</b> <i>egress-vrf</i>   <b>blackhole</b>   <b>reject</b> } [ <i>distance</i> ] [ <b>tag</b> <i>tag-id</i> ]
2.	In configuration mode, configure loopback and physical addresses and acquire interface configuration mode.	<b>interface</b> { <b>vlan</b> <i>vlan-id</i>   <b>loopback</b> <i>loopback-id</i>   <i>interface-name</i> }
3.	In interface configuration mode, configure the IP address for the interface that serves as the BGP speaker.	<b>ip address</b> { <i>ip-address ip-mask</i>   <i>ip-address/prefixLength</i> } [ <b>primary</b>   <b>secondary</b> ]
4.	In configuration mode, specify an AS number for the router and enter BGP Configuration mode.	<b>router bgp</b> <i>as-number</i>
5.	In BGP configuration mode, configure a BGP-specific router ID to override the global router ID.	<b>bgp router-id</b> <i>router-id</i>
6.	In BGP configuration mode, specify the network you want routes imported from and advertised to.	<b>network</b> <i>prefix/length</i> [ <b>route-map</b> <i>name</i> ][ <b>aspath-limit</b> <i>limit</i> ] [ <b>origin</b> <i>code</i> ] [ <b>med</b> <i>value</i> ] [ <b>local-pref</b> <i>value</i> ]
7.	In BGP configuration mode, configure the peer by identifying its IP address and AS.	<b>neighbor</b> <i>ip-address</i> <b>remote-as</b> <i>as-num</i>
8.	In BGP configuration mode, redistribute routes into BGP, optionally specifying a route-map. Supported <i>route-types</i> are connected, static, OSPF, and RIP.	<b>redistribute</b> <i>route-type</i> [ <b>aspath-limit</b> <i>limit</i> ] [ <b>origin</b> <i>code</i> ] [ <b>med</b> <i>value</i> ] [ <b>local-pref</b> <i>value</i> ] [ <b>route-map</b> <i>name</i> ]

## Configuring BGP Neighbor Parameters

Table 35-3 describes the configuring of BGP neighbor parameters.

**Table 35-3 BGP Neighbor Configuration**

Task	Command
In BGP configuration mode, configure the remote AS for the peer.	<b>neighbor</b> <i>ip-address</i> <b>remote-as</b> <i>as-num</i>
In BGP configuration mode, optionally configure EBGP peer routes to not contain this neighbor's AS.	<b>neighbor</b> <i>ip-address</i> <b>ignore-leading-as</b>
In BGP configuration mode, optionally remove private autonomous system (AS) numbers from outbound updates to an external peer.	<b>neighbor</b> { <i>ip-address</i>   <i>groupID</i> } <b>remove-private-as</b>
In BGP configuration mode, optionally modify the minimum interval between the sending of EBGP routing updates.	<b>neighbor</b> <i>ip-address</i> <b>advertisement-interval</b> <i>interval</i>
In BGP configuration mode, optionally configure the conditional advertisement of routes for this neighbor.	<b>neighbor</b> { <i>ip-address</i>   <i>groupID</i> } <b>advertise-map</b> <i>adv-map</i> <b>non-exist-map</b> <i>non-exit-map</i>
In BGP configuration mode, optionally enable the inclusion of confederation information in the AS paths sent to this router's peers.	<b>neighbor</b> { <i>ip-address</i>   <i>groupID</i> } <b>aggregate-confed</b>
In BGP configuration mode, optionally modify the interval between successive update messages for route prefixes that originate in the local AS.	<b>neighbor</b> <i>ip-address</i> <b>as-origination-interval</b> <i>interval</i>
In BGP configuration mode, optionally enable checking to see if the next hop is the peer's address and do not send routes if it is.	<b>neighbor</b> <i>ip-address</i> <b>check-next-hop</b>
In BGP configuration mode, optionally clear all BGP counters for this peer.	<b>neighbor</b> <i>ip-address</i> <b>clear-counters</b>
In BGP configuration mode, optionally configure the specified neighbor as a member of the router's confederation.	<b>neighbor</b> <i>ip-address</i> <b>confed-member</b>
In BGP configuration mode, optionally modify the amount of time between attempts to reestablish a connection to configured peers that are no longer available.	<b>neighbor</b> <i>ip-address</i> <b>connect-retry-interval</b> <i>interval</i>
In BGP configuration mode, optionally force the advertisement of the default route regardless of whether the default route is present in the local routing table.	<b>neighbor</b> <i>ip-address</i> <b>default-originate</b> [ <b>route-map</b> <i>name</i> ]
In BGP configuration mode, explicitly enable a peer that has been administratively disabled. A configured peer is enabled by default.	<b>neighbor</b> <i>ip-address</i> <b>enable</b>
In BGP configuration mode, optionally modify the interval between returning to the idle state and reinitiating a TCP connection for this neighbor.	<b>neighbor</b> <i>ip-address</i> <b>idle-hold-interval</b> <i>interval</i>
In BGP configuration mode, optionally configure EBGP peer routes to not contain this neighbor's AS.	<b>neighbor</b> <i>ip-address</i> <b>ignore-leading-as</b>
In BGP configuration mode, optionally modify the maximum number of Outbound Route Filtering (ORF) entries that will be accepted from this neighbor.	<b>neighbor</b> <i>ip-address</i> <b>maximum-orf</b> <i>num</i>

**Table 35-3 BGP Neighbor Configuration (continued)**


Task	Command
In BGP configuration mode, optionally modify the peak number of prefixes that BGP will accept for installation into the Routing Information Base (RIB).	<b>neighbor</b> <i>ip-address</i> <b>maximum-prefix</b> <i>num</i> [ <b>warning-only</b> ]
In BGP configuration mode, optionally configure BGP to always set the BGP next hop to the EBGP peer's address, overriding third-party next hops.	<b>neighbor</b> { <i>ip-address</i>   <i>groupID</i> } <b>next-hop-peer</b>
In BGP configuration mode, optionally set this neighbor's next hop as the router's own address on advertisement.	<b>neighbor</b> { <i>ip-address</i>   <i>groupID</i> } <b>next-hop-self</b>
In BGP configuration mode, optionally modify the interval between the establishment of a TCP connection and the sending of an OPEN message to open a BGP session.	<b>neighbor</b> <i>ip-address</i> <b>open-delay</b> <i>seconds</i>
In BGP configuration mode, optionally prevent the router from ever trying to open a BGP connection with the specified peer.	<b>neighbor</b> <i>ip-address</i> <b>passive</b>
In BGP configuration mode, optionally create a BGP peer group and add a peer group neighbor.	<b>neighbor</b> <i>groupID</i> <b>peer-group</b> <b>neighbor</b> <i>ip-address</i> <b>peer-group</b> <i>pgname</i>
In BGP configuration mode, optionally configure the peer type of the specified peer or group.	<b>neighbor</b> { <i>ip-prefix/length</i>   <i>groupID</i> } <b>peer-type</b> { <b>ibgp</b>   <b>ebgp</b>   <b>ebgp-confed</b> }
In BGP configuration mode, optionally configure whether updates for prefixes containing the NOPEER community will be accepted by or sent to this neighbor.	<b>neighbor</b> { <i>ip-address</i>   <i>groupID</i> } <b>peering-type</b> { <b>bilateral</b>   <b>unspecified</b> }
In BGP configuration mode, optionally specify a BGP route-map to be used for controlling the import or export of routes to and from the specified peer or group.	<b>neighbor</b> { <i>ip-address</i>   <i>groupID</i> } <b>route-map</b> <i>rm-name</i> { <b>in</b>   <b>out</b> }
In BGP configuration mode, optionally specify that the router will act as a route reflector for this neighbor.	<b>neighbor</b> <i>ip-address</i> <b>route-reflector-client</b>
In BGP configuration mode, optionally modify the interval between the advertisement and subsequent withdrawal of a route.	<b>neighbor</b> <i>ip-address</i> <b>route-withdraw-interval</b> <i>interval</i>
In BGP configuration mode, optionally enable soft-reconfiguration for a peer or peer group.	<b>neighbor</b> { <i>ip-address</i>   <i>groupID</i> } <b>soft-reconfiguration</b>
In BGP configuration mode, optionally modify holdtime and keepalive time values within a BGP peer.	<b>neighbor</b> <i>ip-address</i> <b>timers</b> <i>keepalive-value</i> <i>holdtime-value</i>
In BGP configuration mode, optionally modify the time to live (TTL) value.	<b>neighbor</b> <i>ip-address</i> <b>tll</b> <i>tll-num</i>
In BGP configuration mode, optionally specify the source IP address to be used in all TCP and BGP messages sent to the peer.	<b>neighbor</b> <i>ip-address</i> <b>update-source</b> <i>source-addr</i>

## Configuring Source IP Address Update

Figure 35-6 displays an example of a Source IP address to remote peer configuration. Router 1 configures the loop-back source address 3.3.3.3 on Router 3 as its neighbor. Router 3 configures the loop-back source address 1.1.1.1 on Router 1 as its neighbor. When the outgoing interface on Router 3 for the initial route fails, because the loop-back source address 3.3.3.3 is still operational, the BGP message for this route is not discarded, but rather uses the new route after BGP reconverges.



Figure 35-6 Source IP Address to a Remote Peer



Router 1	Router 3
AS: 10	AS: 10
Source-address: 1.1.1.1	Source-address: 3.3.3.3
Neighbor: 3.3.3.3 remote-as 10	Neighbor: 1.1.1.1 remote-as 10

**Router 1**

```
7100 Chassis(su-config)->router bgp 10
7100 Chassis(su-config-bgp)->bgp router-id 1.1.1.1
7100 Chassis(su-config-bgp)->neighbor 3.3.3.3 remote-as 10
7100 Chassis(su-config-bgp)->neighbor 3.3.3.3 update-source 1.1.1.1
```

**Router 3**

```
7100 Chassis(su-config)->router bgp 10
7100 Chassis(su-config-bgp)->bgp router-id 3.3.3.3
7100 Chassis(su-config-bgp)->neighbor 1.1.1.1 remote-as 10
7100 Chassis(su-config-bgp)->neighbor 1.1.1.1 update-source 3.3.3.3
```

Procedure 35-6 describes how to configure the source IP address to the remote peer.

**Procedure 35-6 Configuring Source IP Address to the Peer Update**


Step	Task	Command(s)
1.	In BGP configuration mode, specify the neighbor this update source IP address will be applied to.	<b>neighbor</b> <i>ip-address</i> <b>remote-as</b> <i>as-num</i>
2.	In BGP configuration mode, specify the update source IP address for this neighbor.	<b>neighbor</b> <i>ip-address</i> <b>update-source</b> <i>source-addr</i>

## Configuring BGP Confederations

The following confederation configuration example presents a confederation of three routers, each belonging to different ASs, with a single EBGP connection to a router outside of the confederation. Router 1 sees Router 2 as belonging to AS (confederation) 100. Router 2 sees Router 1 as belonging to AS 1. Within the confederation, Router 2 - 4, see each other belonging to their respective ASs: 2, 3, and 4.

Figure 35-7 displays the topology for this BGP confederation example.

**Figure 35-7 BGP Confederation Example Topology**



### Router 1

```
Router 1 (rw) -> configure
Router 1 (rw-config) -> interface vlan 1
Router 1 (rw-config-intf-vlan.0.1) -> ip address 200.10.1.1 255.255.255.0
Router 1 (rw-config-intf-vlan.0.1) -> no shutdown
Router 1 (rw-config-intf-vlan.0.1) -> exit

Router 1 (rw) -> configure
Router 1 (rw-config) -> router bgp 1
Router 1 (su-config-bgp) -> bgp router-id 1.1.1.1
Router 1 (su-config-bgp) -> neighbor 200.10.1.2 remote-as 100
Router 1 (su-config-bgp) -> redistribute static
Router 1 (su-config-bgp) ->
```

## Router 2

```
Router 2 (rw)->configure
Router 2 (rw-config)->interface vlan 1
Router 2 (rw-config-intf-vlan.0.1)->ip address 200.10.1.2 255.255.255.0
Router 2 (rw-config-intf-vlan.0.1)->no shutdown
Router 2 (rw-config-intf-vlan.0.1)->exit

Router 2 (rw-config)->interface vlan 2
Router 2 (rw-config-intf-vlan.0.2)->ip address 200.10.2.1 255.255.255.0
Router 2 (rw-config-intf-vlan.0.2)->no shutdown
Router 2 (rw-config-intf-vlan.0.2)->exit

Router 2 (rw-config)->interface vlan 4
Router 2 (rw-config-intf-vlan.0.4)->ip address 200.10.4.1 255.255.255.0
Router 2 (rw-config-intf-vlan.0.4)->no shutdown
Router 2 (rw-config-intf-vlan.0.4)->exit

Router 2 (rw-config)->router bgp 2
Router 2 (su-config-bgp)->bgp router-id 2.2.2.2
Router 2 (su-config-bgp)->neighbor 200.10.1.1 remote-as 1
Router 2 (su-config-bgp)->neighbor 200.10.2.2 remote-as 3
Router 2 (su-config-bgp)->neighbor 200.10.4.2 remote-as 4
Router 2 (su-config-bgp)->bgp confederation-id 100
Router 2 (su-config-bgp)->neighbor 200.10.2.2 confed-member
Router 2 (su-config-bgp)->neighbor 200.10.4.2 confed-member
Router 2 (su-config-bgp)->redistribute static
Router 2 (su-config-bgp)->
```

## Router 3

```
Router 3 (rw)->configure
Router 3 (rw-config)->interface vlan 2
Router 3 (rw-config-intf-vlan.0.2)->ip address 200.10.2.2 255.255.255.0
Router 3 (rw-config-intf-vlan.0.2)->no shutdown
Router 3 (rw-config-intf-vlan.0.2)->exit

Router 3 (rw-config)->interface vlan 3
Router 3 (rw-config-intf-vlan.0.3)->ip address 200.10.3.1 255.255.255.0
Router 3 (rw-config-intf-vlan.0.3)->no shutdown
Router 3 (rw-config-intf-vlan.0.3)->exit

Router 3 (rw-config)->router bgp 3
Router 3 (su-config-bgp)->bgp router-id 3.3.3.3
Router 3 (su-config-bgp)->neighbor 200.10.2.1 remote-as 2
Router 3 (su-config-bgp)->neighbor 200.10.3.2 remote-as 4
```

```

Router 3 (su-config-bgp)->bgp confederation-id 100
Router 3 (su-config-bgp)->neighbor 200.10.2.1 confed-member
Router 3 (su-config-bgp)->neighbor 200.10.3.2 confed-member
Router 3 (su-config-bgp)->redistribute static
Router 3 (su-config-bgp)->

```

## Router 4

```

Router 4 (rw)->configure
Router 4 (rw-config)->interface vlan 3
Router 4 (rw-config-intf-vlan.0.3)->ip address 200.10.3.2 255.255.255.0
Router 4 (rw-config-intf-vlan.0.3)->no shutdown
Router 4 (rw-config-intf-vlan.0.3)->exit

```

```

Router 4 (rw-config)->interface vlan 4
Router 4 (rw-config-intf-vlan.0.4)->ip address 200.10.4.2 255.255.255.0
Router 4 (rw-config-intf-vlan.0.4)->no shutdown
Router 4 (rw-config-intf-vlan.0.4)->exit

```

```

Router 4 (rw-config)->router bgp 4
Router 4 (su-config-bgp)->bgp router-id 4.4.4.4
Router 4 (su-config-bgp)->neighbor 200.10.3.1 remote-as 3
Router 4 (su-config-bgp)->neighbor 200.10.4.1 remote-as 2
Router 4 (su-config-bgp)->bgp confederation-id 100
Router 4 (su-config-bgp)->neighbor 200.10.3.1 confed-member
Router 4 (su-config-bgp)->neighbor 200.10.4.1 confed-member
Router 4 (su-config-bgp)->redistribute static
Router 4 (su-config-bgp)->

```

[Procedure 35-7](#) describes how to configure BGP confederations.

### Procedure 35-7 Configuring BGP Confederation


Step	Task	Command(s)
1.	In BGP configuration mode, specify the confederation this BGP router belongs to.	<b>bgp confederation identifier</b> <i>confed-id</i>
2.	In BGP configuration mode, configure the specified neighbor as a member of the router's confederation.	<b>neighbor</b> <i>ip-address</i> <b>confed-member</b>
3.	In BGP configuration mode, optionally enable the inclusion of confederation information in the AS paths sent to this router's peers.	<b>neighbor</b> <i>{ip-address   peer-group}</i> <b>aggregate-confed</b>
4.	In BGP configuration mode, optionally enable BGP to drop AS-Paths with erroneous confederation segments.	<b>bgp strict-confeds</b>

## Configuring Route Reflection

The following route reflection example configures a single route reflector with two clients all of which are members of AS 1. Router 1 is the route reflector. Routers 2 and 3 are route reflector clients. Router 2 has an EBGP connection to Router 4 of AS 2. Router 3 has an EBGP connection to Router 5 of AS 3. Routers 2 and 3 advertise all their routes to Router 1. Router 1 advertises routes learned from Router 2 to Router 3 and routes learned from Router 3 to Router 2.

Figure 35-8 displays the BGP route reflection topology for this example.

**Figure 35-8 BGP Route Reflection Example Topology**



### Router 1

```
Router 1 (rw) -> configure
Router 1 (rw-config) -> interface vlan 1
Router 1 (rw-config-intf-vlan.0.1) -> ip address 200.10.1.1 255.255.255.0
Router 1 (rw-config-intf-vlan.0.1) -> no shutdown
Router 1 (rw-config-intf-vlan.0.1) -> exit

Router 1 (rw) -> configure
Router 1 (rw-config) -> interface vlan 2
Router 1 (rw-config-intf-vlan.0.2) -> ip address 200.10.2.1 255.255.255.0
Router 1 (rw-config-intf-vlan.0.2) -> no shutdown
```

```
Router 1 (rw-config-intf-vlan.0.2)->exit
```

```
Router 1 (rw)->configure
```

```
Router 1 (rw-config)->router bgp 1
```

```
Router 1 (su-config-bgp)->bgp router-id 1.1.1.1
```

```
Router 1 (su-config-bgp)->bgp cluster-id 1.1.1.1
```

```
Router 1 (su-config-bgp)->neighbor 200.10.1.2 remote-as 1
```

```
Router 1 (su-config-bgp)->neighbor 200.10.1.2 route-reflector-client
```

```
Router 1 (su-config-bgp)->neighbor 200.10.2.2 remote-as 1
```

```
Router 1 (su-config-bgp)->neighbor 200.10.2.2 route-reflector-client
```

```
Router 1 (su-config-bgp)->redistribute static
```

```
Router 1 (su-config-bgp)->
```

## Router 2

```
Router 2 (rw)->configure
```

```
Router 2 (rw-config)->interface vlan 1
```

```
Router 2 (rw-config-intf-vlan.0.1)->ip address 200.10.1.2 255.255.255.0
```

```
Router 2 (rw-config-intf-vlan.0.1)->no shutdown
```

```
Router 2 (rw-config-intf-vlan.0.1)->exit
```

```
Router 2 (rw)->configure
```

```
Router 2 (rw-config)->interface vlan 3
```

```
Router 2 (rw-config-intf-vlan.0.2)->ip address 200.10.3.1 255.255.255.0
```

```
Router 2 (rw-config-intf-vlan.0.2)->no shutdown
```

```
Router 2 (rw-config-intf-vlan.0.2)->exit
```

```
Router 2 (rw)->configure
```

```
Router 2 (rw-config)->router bgp 1
```

```
Router 2 (su-config-bgp)->bgp router-id 2.2.2.2
```

```
Router 2 (su-config-bgp)->neighbor 200.10.1.1 remote-as 1
```

```
Router 2 (su-config-bgp)->neighbor 200.10.3.2 remote-as 2
```

```
Router 2 (su-config-bgp)->redistribute static
```

```
Router 2 (su-config-bgp)->
```

## Router 3

```
Router 3 (rw)->configure
```

```
Router 3 (rw-config)->interface vlan 2
```

```
Router 3 (rw-config-intf-vlan.0.2)->ip address 200.10.2.2 255.255.255.0
```

```
Router 3 (rw-config-intf-vlan.0.2)->no shutdown
```

```
Router 3 (rw-config-intf-vlan.0.2)->exit
```

```
Router 3 (rw)->configure
```

```
Router 3 (rw-config)->interface vlan 4
```

```
Router 3 (rw-config-intf-vlan.0.4)->ip address 200.10.4.1 255.255.255.0
```

```
Router 3 (rw-config-intf-vlan.0.4)->no shutdown
```

```
Router 3 (rw-config-intf-vlan.0.4)->exit
```

```

Router 3 (rw)->configure
Router 3 (rw-config)->router bgp 1
Router 3 (su-config-bgp)->bgp router-id 3.3.3.3
Router 3 (su-config-bgp)->neighbor 200.10.2.1 remote-as 1
Router 3 (su-config-bgp)->neighbor 200.10.4.2 remote-as 3
Router 3 (su-config-bgp)->redistribute static
Router 3 (su-config-bgp)->

```

### Router 4

```

Router 4 (rw)->configure
Router 4 (rw-config)->interface vlan 3
Router 4 (rw-config-intf-vlan.0.3)->ip address 200.10.3.2 255.255.255.0
Router 4 (rw-config-intf-vlan.0.3)->no shutdown
Router 4 (rw-config-intf-vlan.0.3)->exit

```

```

Router 4 (rw)->configure
Router 4 (rw-config)->router bgp 2
Router 4 (su-config-bgp)->bgp router-id 4.4.4.4
Router 4 (su-config-bgp)->neighbor 200.10.3.1 remote-as 1
Router 4 (su-config-bgp)->redistribute static
Router 4 (su-config-bgp)->

```

### Router 5

```

Router 5 (rw)->configure
Router 5 (rw-config)->interface vlan 4
Router 5 (rw-config-intf-vlan.0.4)->ip address 200.10.4.2 255.255.255.0
Router 5 (rw-config-intf-vlan.0.4)->no shutdown
Router 5 (rw-config-intf-vlan.0.4)->exit

```

```

Router 5 (rw)->configure
Router 5 (rw-config)->router bgp 3
Router 5 (su-config-bgp)->bgp router-id 5.5.5.5
Router 5 (su-config-bgp)->neighbor 200.10.4.1 remote-as 1
Router 5 (su-config-bgp)->redistribute static
Router 5 (su-config-bgp)->

```

[Procedure 35-8](#) describes how to configure BGP route reflection.

### Procedure 35-8 Configuring BGP Route Reflection

Step	Task	Command(s)
1.	In BGP configuration mode, specify that the router will act as a route reflector for the specified neighbor.	<b>neighbor <i>ip-address</i> route-reflector-client</b>

**Procedure 35-8 Configuring BGP Route Reflection (continued)**

Step	Task	Command(s)
2.	In BGP configuration mode, specify the route reflection cluster ID for the cluster the route reflector belongs to. This value defaults to the route reflector router ID if only a single route reflector is configured for the cluster, otherwise a cluster ID must be specified.	<b>bgp cluster-id</b> <i>router-id</i>

## Configuring Outbound Route Filtering (ORF)

[Table 35-4](#) describes how to configure BGP Outbound Route Filtering.

**Table 35-4 Configuring BGP Outbound Route Filtering**

Task	Command(s)
In BGP configuration mode, optionally specify whether the ORF capability for community filtering is to be sent to the BGP peer, received from the BGP peer, or both.	<b>bgp orf {ipv4   ipv6} unicast comm-filter {send   receive   both}</b>
In BGP configuration mode, optionally specify whether the ORF capability for extended community filtering is to be sent to the BGP peer, received from the BGP peer, or both.	<b>bgp orf {ipv4   ipv6} unicast extcomm-filter {send   receive   both}</b>
In BGP configuration mode, optionally specify whether the Outbound Route Filtering (ORF) capability for prefix filtering is to be sent to the BGP peer, received from the BGP peer, or both.	<b>bgp orf {ipv4   ipv6} unicast prefix-filter {send   receive   both}</b>

## Configuring Conditional Advertisement


The conditional advertisement feature allows a service provider to advertise certain routes to a preferred subnet under normal operational conditions, while maintaining the ability to move its traffic to an alternative subnet should its preferred routes fail.

The following example configures two static routes on Router 1. One will be advertised to Router 2 under normal operational conditions, the other will only be advertised if the first route no longer exists. A prefix list is created for each static route destination prefix. The destination prefix to be advertised under normal operational conditions is assigned to the non-exist-map route-map. The destination prefix to be advertised only if the advertised prefix fails is assigned to the advertise route-map.

[Figure 35-9](#) on page 35-43 presents the conditional advertisement example topology.

**Figure 35-9 BGP Conditional Advertisement Example Topology**





### For Router 1

- Configure two static routes with next hops to **192.168.13.112** on **VLAN 13** to destination prefixes **1.0.0.0/8** and **2.0.0.0/8**
- Configure the non-exist map prefix-list named **non-exist-list1** and assign it to BGP route-map **non-exist-map1**, specifying prefix **2.0.0.0/8** as the non-exist map prefix
- Configure the advertise map prefix-list named **adv-list1** and assign it to BGP route-map **adv-map1**, specifying prefix **1.0.0.0/8** as the advertised prefix
- Configure the router for **AS 10** with a router ID of **1.1.1.1**
- Configure IP address **192.168.12.112** as the peer
- Configure the BGP advertise map for neighbor **192.168.12.112** and assign **adv-map1** as the advertise map route-map and **non-exist-map1** as the non-exist map route-map

```

Router 1 (su)->configure
Router 1 (su-config)->ip route 1.0.0.0/8 192.168.13.112 interface vlan.0.13
Router 1 (su-config)->ip route 2.0.0.0/8 192.168.13.112 interface vlan.0.13
Router 1 (su-config)->ip prefix-list adv-list1 permit seq 1 1.0.0.0/8
Router 1 (su-config)->route-map bgp adv-map1
Router 1 (su-config-route-map-bgp)->match prefix-list adv-list1
Router 1 (su-config-route-map-bgp)->exit
Router 1 (su-config)->ip prefix-list non-exist-list1 permit seq 1 2.0.0.0/8
Router 1 (su-config)->route-map bgp non-exist-map1
Router 1 (su-config-route-map-bgp)->match prefix-list non-exist-list1
Router 1 (su-config-route-map-bgp)->exit
Router 1 (su-config)->router bgp 10
Router 1 (su-config-bgp)->bgp router-id 1.1.1.1

```

```

Router 1(su-config-bgp)->neighbor 192.168.12.112 remote-as 10
Router 1(su-config-bgp)->neighbor 192.168.12.112 advertise-map adv-map1
non-exist-map non-exist-map1
Router 1(su-config-bgp)->

```

### For Router 2:

- Configure a static route with the next hop **192.168.12.111** on **VLAN 12** to destination prefix **192.168.13.0/24**
- Configure Router 2 for AS **10** with a router ID of **2.2.2.2**
- Configure IP address **192.168.12.111** as the peer

```

Router 2(su)->configure
Router 2(su-config)->ip route 192.168.13.0/24 192.168.12.111 interface vlan.0.12
Router 2(su-config)->router bgp 10
Router 2(su-config-bgp)->bgp router-id 2.2.2.2
Router 2(su-config-bgp)->neighbor 192.168.12.111 remote-as 10
Router 2(su-config-bgp)->Verifying the Configuration

```

If advertised map configuration is not applied, the Router 1 advertised routes output for 192.168.12.112 would display both configured static routes as being advertised to Router 2:

```

Router 1(su)->show ip bgp peer 192.168.12.112 advertised-routes
Route status codes: adv - advertised, sup - suppressed, pw - pending w/drawal, wd
- w/drawn
Route aggregation codes:
1 - Route is not aggregating or aggregated
2 - Route is aggregating
3 - Route is unsuppressed aggregated
4 - Route is suppressed aggregated

```

Stat	Aggr	Network	Next Hop	Rib	MED	Local-Pref	Origin	AS	Path
adv	1	1.0.0.0	192.168.13.112	U	0	100	Inc		
adv	1	2.0.0.0	192.168.13.112	U	0	100	Inc		

With the advertised map configuration applied, only network 2.0.0.0 displays:

```

Router 1(su)->show ip bgp peer 192.168.12.112 advertised-routes
....
Stat Aggr  Network          Next Hop          Rib MED Local-Pref Origin AS Path
adv  1          2.0.0.0          192.168.13.112  U  0      100      Inc

```

Should static route 2.0.0.0/8 be withdrawn, network 1.0.0.0 is advertised and network 2.0.0.0 displays as withdrawn:

```

Router 1(su-config)->no ip route 2.0.0.0/8 192.168.13.112 interface vlan.0.13 1
Router 1(su-config)->show ip bgp peer 192.168.12.112 advertised-routes
....
Stat Aggr  Network          Next Hop          Rib MED Local-Pref Origin AS Path
adv  1          1.0.0.0          192.168.13.112  U  0      100      Inc
wd   1          2.0.0.0          192.168.13.112  U  0      100      Inc

```

[Procedure 35-9](#) describes how to configure BGP conditional route advertisement.

**Procedure 35-9 Configuring BGP Conditional Route Advertisement**

Step	Task	Command(s)
1.	In router configuration mode, configure a prefix list for both the advertise and non-exist prefixes to be matched in the appropriate route-maps.	<b>ip prefix-list</b> <i>name</i> [ <b>seq</b> <i>seq-value</i> ] <b>{deny   permit}</b> <i>{prefix/length}</i> [ <i>source-address</i> ] [ <i>next-hop</i> ] [ <b>ge</b> <i>length</i> ] [ <b>le</b> <i>length</i> ] [ <b>niri</b> ]
2.	In router configuration mode, create both an advertise and non-exist route-map.	<b>route-map</b> <b>bgp</b> <i>name</i> [ <b>permit   deny</b> ] [ <i>sequence-number</i> ]
3.	In BGP route-map configuration mode, add match statements containing the configured prefix lists to the appropriate BGP route-map.	<b>match prefix-list</b> <i>prefix-list</i>
4.	In BGP configuration mode, specify the advertise and non-exist maps to be applied to this neighbor.	<b>neighbor</b> <i>{ip-address   groupID}</i> <b>advertise-map</b> <i>adv-map</i> <b>non-exist-map</b> <i>non-exist-map</i>

## Configuring BGP Soft Reset

Table 35-5 describes how to configure BGP soft reset.

**Table 35-5 Configuring BGP Soft Reset**


Task	Command(s)
In BGP configuration mode, optionally enable BGP soft reconfiguration for a peer or peer group.	<b>neighbor</b> <i>{ip-address   groupID}</i> <b>soft-reconfiguration</b>
In BGP configuration mode, optionally disable the automatic sending of route-refresh messages on inbound policy changes. Automatic route refresh is enabled by default. If soft reconfiguration is enabled, route refreshes are not sent.	<b>no bgp automatic-route-refresh</b>
In any command mode, optionally tear down one or all BGP connections. Optionally specify the <b>soft</b> option to perform a route refresh.	<b>clear ip bgp</b> <i>{peer-address   *}</i> [ <b>soft</b> ]

## Configuring Flap Dampening

This section presents a flap dampening configuration example. The example configures a static route on Router 1 that is redistributed to Router 2. On Router 2 a flap table is configured and applied to an inbound route-map that will monitor the link between Router 2 and Router 1. Using show output, the example tracks the changes in route penalty over two flaps of the route that result in the route being suppressed. After waiting for the 5 minute hold-time to expire, the show output displays the route as unsuppressed.

Figure 35-10 on page 35-46 presents the route flap dampening example configuration.

**Figure 35-10 Route Flap Dampening Example Configuration**



Router 1	Router 2
<b>AS:</b> 10	<b>AS:</b> 5
<b>Static Route:</b>	<b>Neighbor:</b> 192.168.12.111 remote-as 10
Destination Prefix: 1.0.0.0/8	<b>Flap table:</b>
Next Hop: 192.168.13.112	Name: flap1
<b>Neighbor:</b> 192.168.12.113 remote-as 5	Cutoff: 150
<b>Redistribute:</b> static	Reuse: 75
	Hold-time: 300 seconds
	<b>Route-map:</b> map1

## Router 1

Router 1 is configured by:

- Configuring a static route on the 192.168.13.0/24 subnet with a next hop of **192.168.13.112** and a destination address of **1.0.0.0/8** on Router 1 (AS 10)
- Setting the router ID to **1.2.3.4**
- Configuring Router 1 with a BGP neighbor **192.168.12.112** (AS 5)
- Redistributing the static route to Router 2

```
7100 Chassis(su)->configure
7100 Chassis(su-config)->ip route 1.0.0.0/8 192.168.13.112 interface vlan.0.13
7100 Chassis(su-config)->bgp router 10
7100 Chassis(su-config-bgp)->bgp router-id 1.2.3.4
7100 Chassis(su-config-bgp)->neighbor 192.168.12.112 remote-as 5
7100 Chassis(su-config-bgp)->redistribute static
```

## Router 2

Router 1 is configured by:

- Configuring the **flap1** flap table on Router 2 with:
  - **cutoff** set to **150**
  - **reuse** set to **75**
  - **hold-time** set to **300**
- Applying the **flap1** flap table to the **map1** BGP route-map
- Configuring Router 2 with a BGP neighbor **192.168.12.111** (AS 10)
- Applying the **map1** inbound route-map to neighbor **192.168.12.111**
- Clear the BGP session to initiate a readvertisement of the route from Router 1 to Router 2
- Display the route and route flap dampening details

```
7100 Chassis(su)->configure
7100 Chassis(su-config)->dampen-flap flap1
S Chassis(su-config-dampen-flap)->cutoff 150
S Chassis(su-config-dampen-flap)->reuse 75
S Chassis(su-config-dampen-flap)->hold-time 300
S Chassis(su-config-dampen-flap)->exit
7100 Chassis(su-config)->route-map bgp map1 permit 10
7100 Chassis(su-config-route-map-bgp)->set flap-table flap1
7100 Chassis(su-config-route-map-bgp)->exit
S Chassis(su-config)->router bgp 5
7100 Chassis(su-config-bgp)->bgp router-id 2.3.4.5
7100 Chassis(su-config-bgp)->neighbor 192.168.12.111 remote-as 10
7100 Chassis(su-config-bgp)->neighbor 192.168.12.111 route-map map1 in
7100 Chassis(su-config-bgp)->exit
7100 Chassis(su-config)->
```

The following displays route dampening statistics for network 1.0.0.0/8 prior to any instability:

```
7100 Chassis(su-config)->show ip bgp 1.0.0.0/8 detail
Route status codes: > - active
```

Network	Next Hop	Rib MED	Local-Pref	Origin	AS Path
> 1.0.0.0/8	192.168.12.111	U	0	100	Inc 10

Community attributes in route:

Extended Community attributes in route:

Route Flap Dampening configuration file name: flap1

Is route suppressed? No

Flap penalty: 0, Flap Count 0, Flap time remaining 0 seconds

```
7100 Chassis(su-config)->
```

If you were to enter the show command immediately after the first flap, the route flap dampening statistics would show:

```
7100 Chassis(su-config)->show ip bgp 1.0.0.0/8 detail
Route status codes: > - active
...
Route Flap Dampening configuration file name: flap1
Is route suppressed? No
Flap penalty: 99, Flap Count 1, Flap time remaining 0 seconds
```

The route remains unsuppressed because the route penalty has not exceeded the cutoff value. The flap penalty has decayed by 1 point from 100 to 99 since the route flap occurred. The flap count is 1 and because the route has not yet been suppressed, flap time is set to 0 seconds.

Entering the show command immediately after the second route flap, the route flap dampening statistics would show:

```
7100 Chassis(su-config)->show ip bgp 1.0.0.0/8 detail
Route status codes: > - active
...
Route Flap Dampening configuration file name: flap1
Is route suppressed? Yes
Flap penalty: 185, Flap Count 2, Flap time remaining 295 seconds
```

Because the flap penalty exceeds the cutoff setting, route flap dampening has suppressed the route. Flap time now displays the amount of seconds remaining before the hold-time is reached.

After waiting the 5 minute hold-time, the display shows that the route has been unsuppressed. If the hold-time was set to the default value of 900 seconds, and no other router flaps occurred for this route, the route penalty would have decayed to the point that the route would have been unsuppressed when the route penalty reached the reuse setting of 75.

```
7100 Chassis(su-config)->show ip bgp 1.0.0.0/8 detail
Route status codes: > - active
...
Route Flap Dampening configuration file name: flap1
Is route suppressed? No
Flap penalty: 86, Flap Count 2, Flap time remaining 0 seconds
```

[Procedure 35-10](#) describes how to configure the BGP flap dampening.

### Procedure 35-10 Configuring BGP Flap Dampening

Step	Task	Command(s)
1.	In router configuration mode, enter the dampen flap command mode for the named dampen flap table.	<b>dampen-flap</b> <i>name</i>
2.	In BGP flap dampening configuration mode, optionally modify the route suppression threshold.	<b>cutoff</b> <i>threshold</i>
3.	In BGP flap dampening configuration mode, optionally modify the time in seconds after which a reachable route's penalty value decays to half of its current value.	<b>half-life-reach</b> <i>seconds</i>

**Procedure 35-10 Configuring BGP Flap Dampening**

Step	Task	Command(s)
4.	In BGP flap dampening configuration mode, optionally modify the time in seconds after which an unreachable route's penalty value decays to half its current value.	<b>half-life-unreach</b> <i>seconds</i>
5.	In BGP flap dampening configuration mode, optionally modify the maximum time a route can be suppressed.	<b>hold-time</b> <i>seconds</i>
6.	In BGP flap dampening configuration mode, optionally modify the decay memory limit for reachable routes.	<b>memory-limit-reach</b> <i>seconds</i>
7.	In BGP flap dampening configuration mode, optionally modify the decay memory limit for unreachable routes.	<b>memory-limit-unreach</b> <i>seconds</i>
8.	In BGP flap dampening configuration mode, optionally modify the route penalty value below which a suppressed route is reused.	<b>reuse</b> <i>value</i>

**Configuring Graceful Restart**

[Procedure 35-11](#) describes how to configure graceful restart.

**Procedure 35-11 Configuring Graceful Restart**

Step	Task	Command(s)
1.	In BGP configuration mode, optionally enable graceful restart on the router.	<b>bgp graceful-restart</b>
2.	In BGP configuration mode, optionally modify the time to defer route selection after graceful restarting.	<b>bgp restart-defer</b> <i>time-seconds</i>
3.	In BGP configuration mode, optionally modify the maximum time to wait for a graceful restart capable peer to come back after a restart.	<b>bgp restart-time</b> <i>time-seconds</i>
4.	In BGP configuration mode, optionally modify the estimated time advertised to peers in the OPEN message for the session to be reestablished after a graceful restart.	<b>bgp restart-timeout</b> <i>time-seconds</i>
5.	In BGP configuration mode, optionally modify the maximum time following a restart before removing stale routes from the peer.	<b>bgp stale-path-time</b> <i>time-seconds</i>

## BGP Monitoring and Clearing

Table 35-6 describes how to monitor and clear BGP configuration.

**Table 35-6 Monitoring and Clearing BGP Configuration**

Task	Command(s)
In any command mode, optionally reset BGP peering sessions and optionally send route refresh requests.	<b>clear ip bgp</b> {peer   *} [soft]
In any command mode, optionally clear all route-flap statistics and state for the specified route-prefix.	<b>clear ip bgp flap-all-stats</b> ip-prefix/length
In any command mode, optionally clear the route-flap count for the specified route-prefix.	<b>clear ip bgp flap-count</b> ip-prefix/length
In any command mode, optionally display information about BGP routes installed in the BGP routing information base (RIB).	<b>show ip bgp</b> [ip-address] [ip-prefix/mask] [longer-prefixes] [detail] [peer ip-addr {all-received-routes   received-routes   advertised-routes}]
In any command mode, optionally display dampened routes information.	<b>show ip bgp dampened-routes</b>
In any command mode, optionally display information for BGP groups.	<b>show ip bgp groups</b>
In any command mode, optionally display information about the state of BGP's IPv4 peering sessions.	<b>show ip bgp neighbors</b> [ip-address]
In any command mode, optionally display a summary of the BGP configuration.	<b>show ip bgp summary</b>

## Terms and Definitions

Table 35-7 lists terms and definitions used in this BGP configuration discussion.

**Table 35-7 BGP Terms and Definitions**

Term	Definition
4-octet AS numbers	A BGP extension, defined in RFC 4893, that allows for the encoding of 4-octet AS numbers.
aggregation	A BGP feature that provides for the aggregating of one or more specific routes into a single aggregate route, if a more specific route of the aggregate route exists in the BGP routing table.
AS	A set of routers under a single administration referred to as an autonomous system.
AS path	A BGP path attribute, used for loop detection, that provides a list of the AS numbers the route traverses.
BGP	The Border Gateway Protocol (BGP), defined in RFC 4271, that is the standard Exterior Gateway Protocol (EGP) for routing between administrative domains.
conditional advertisement	Conditional BGP announcements that are sent in addition to normal announcements, when a route specified in the configured advertise map does not exist in the configured non-exist map



**Table 35-7 BGP Terms and Definitions (continued)**

Term	Definition
confederation	An extension to BGP, defined in RFC 3065, which may be used to create a confederation of autonomous systems that is represented as a single autonomous system to BGP peers external to the confederation, thereby removing the "full mesh" requirement, aiding in policy administration, and reducing the management complexity of maintaining a large autonomous system.
graceful restart	A capability, defined in RFC 4724, that provides for the continued processing of the data-forwarding plane of a router should the control plane fail.
MED	The Multi-Exit Discriminator attribute used by external neighbors in the selection of the preferred path into an autonomous system (AS) that has multiple entry points.
multi-protocol BGP extensions	A set of BGP extensions, defined in RFC 2858, that enable BGP to carry routing information for multiple Network Layer protocols such as IPv6 and IPX.
outbound route filtering	A BGP capability, defined in RFC 5291, that allows a BGP speaker to send to its BGP peer a set of Outbound Route Filters (ORFs), which the peer applies in addition to its locally configured outbound filters (if any), to constrain its outbound routing updates to the speaker
peer group	A group of neighbors that share the same BGP attributes.
peer-group	A BGP capability that provides for the grouping of peers for purposes of policy, such that the group policy takes precedence for route export, and the peer configured policy takes precedence for route import.
route flap dampening	A BGP capability, defined in RFC 2439, that treats routes that are being announced and withdrawn at a rapid rate as unreachable, based upon a route penalty for each route withdrawal, and reachable again based upon a configurable decay over time of that route penalty.
route reflector	A BGP capability, defined in RFC 4456, that allows a BGP speaker (route reflector) to advertise IBGP learned routes to certain IBGP peers, relieving the scaling issue associated with a fully meshed AS.
route refresh	A BGP capability, defined in RFC 2918, which allows for the dynamic exchange of route refresh requests between BGP speakers and the subsequent re-advertisement of the respective Adj-RIB-Out.
soft reconfiguration	A BGP capability that speeds up the route installation process when an inbound policy change occurs by keeping a local copy of the routes for the specified peer or group.
TCP	Transmission Control Protocol.
TCP/MD5 authentication	A TCP extension to BGP security that defines, in RFC 2385, a TCP option for carrying an MD5 digest in a TCP segment and acts like a signature for that segment, incorporating information known only to the connection end points.

## Virtual Router Redundancy Protocol (VRRP) Configuration

This document describes the Virtual Router Redundancy Protocol (VRRP) feature and its configuration on Extreme Networks 7100-Series devices.

For information about...	Refer to page...
<a href="#">Using VRRP in Your Network</a>	36-1
<a href="#">Implementing VRRP in Your Network</a>	36-2
<a href="#">VRRP Overview</a>	36-2
<a href="#">Configuring VRRP</a>	36-7
<a href="#">VRRP Configuration Examples</a>	36-9
<a href="#">Terms and Definitions</a>	36-13

### Using VRRP in Your Network

Virtual Router Redundancy Protocol (VRRP) is an election protocol capable of dynamically assigning responsibility for a virtual router to one of the VRRP routers on a LAN. A virtual router is an abstract object managed by VRRP that acts as a default router for hosts on a shared LAN. It consists of a Virtual Router Identifier (VRID) and a set of associated IP addresses across a common LAN that define virtual router members. A VRRP router is a router with the VRRP protocol running on it. A VRRP router may participate in and backup one or more virtual routers.

VRRP specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. The elected VRRP router is called the master. The router master controls the IP addresses associated with a virtual router. The master forwards packets sent to these IP addresses. The VRRP election process provides dynamic fail over of forwarding responsibility to another VRRP router should the current master become unavailable. This allows any of the virtual router IP addresses on the LAN to be used as the default first hop router by end-hosts. In this way, VRRP provides a higher availability default path without requiring configuration of dynamic routing or router discovery protocols on every end-host.



**Note:** When configuring an IPv6 VRRP link local address, all link local addresses must match on all routers running the same VRRP instance in a LAN segment. Only one link local address on a VRRP instance will be active at any given time.

Statically configured default routes can represent a single point of failure that can result in a catastrophic event, isolating all end-hosts that are unable to detect any alternate available path. VRRP is designed to eliminate the single point of failure inherent in the static default routed environment.

A critical-IP address defines an interface that will prevent the master router from functioning properly if the interface were to fail. When a critical IP interface goes down, its operational priority can be set to decrement to a value lower than the priority set for the backup router. In this case, the backup router becomes the master.

Fabric route mode can be enabled on the VRRP router allowing a VRRP instance in the backup state to forward IPv4 and IPv6 packets destined for the VRRP MAC address

## Implementing VRRP in Your Network

To implement VRRP:

- Create a virtual router instance
- Configure all VRRP IP addresses associated with this virtual router
- Optionally change the VRRP router priority for this virtual router
- Optionally change the advertise interval for this virtual router
- Optionally set the VRID state interface down transition to interface up delay
- Configure a critical-IP interface, the failure of which will decrement the operational priority of the router causing the backup router to take over as master.
- Optionally configure this virtual router for VRRP authentication (Only applies to VRID's that have been created as version 2 of the protocol)
- Optionally enable accept-mode for this virtual router allowing the master for this virtual router to accept IP packets for the configured associated IP address list
- Optionally enable fabric route mode on the VRRP router.
- Optionally change the master preemption setting for this VRRP router
- Verify configuration and statistics using the VRRP display command

## VRRP Overview

This section provides an overview of VRRP configuration.

### Basic VRRP Topology

**Figure 36-1 A Basic VRRP Topology**




Figure 36-1 shows a basic VRRP topology with a single virtual router. Routers R1 and R2 are both configured with one virtual router (VRID 1). Router R1 serves as the master and Router R2 serves as the backup. The hosts are configured to use 172.111.1.1/16 as the default route.

If Router R1 should become unavailable, Router R2 would take over virtual router VRID 1 and its associated IP addresses. Packets sent to 172.111.1.1/16 would go to Router R2. When Router R1 comes up again, it would take over as master, and Router R2 would revert to backup.

## VRRP Virtual Router Creation

Each virtual router has its own instance. Create a VRRP virtual router instance using the **vrrp create** command in interface configuration command mode specifying the VRID for this instance. The virtual router instance must be created on a routing interface before any other VRRP settings can be configured.

## VRRP Master Election

After the virtual router instance has been created, assign the IP addresses associated with this virtual router using the **vrrp address** command in interface configuration command mode, specifying the IP address and the VRID this address is to be associated with. A virtual router IP address can be either an address configured on the routing interface or an address that falls within the range of any networks configured on this routing interface.

If the virtual router IP address is the same as the routing interface (VLAN) address owned by a VRRP router, then the router owning the address becomes the master. The master sends an advertisement to all other VRRP routers declaring its status and assumes responsibility for forwarding packets associated with its VRID.

If the virtual router IP address is not owned by any of the VRRP routers, then the routers compare their priorities and the higher priority owner becomes the master. VRRP router priority is set using the **vrrp priority** command in interface configuration command mode. If priority values are the same, then the VRRP router with the highest IP address is selected master.

VRRP advertisements are sent by the master router to other routers participating in the VRRP master selection process, informing them of its configured values. Once the master is selected, then advertisements are sent every advertising interval to let other VRRP routers in this VRID know the router is still acting as master of the VRID. All routers with the same VRID should be configured with the same advertisement interval. Use the **vrrp advertise-interval**, in interface configuration command mode, to change the advertise-interval for this VRID.

## Configuring a VRRP Critical-IP Address

A critical-IP address defines an interface that will prevent the master router from functioning properly if the interface were to fail. A critical-IP address is typically an internet facing interface, but can be any IP address that does not include the VRRP configured interface between hosts and a VRRP master or backup first-hop router. An IP address of an interface connecting a master router to a router configured for internet access would be considered a critical-IP address for VRRP routing. Critical-IP addresses can be both local or remote.

Use the **vrrp critical-ip** command in interface configuration command mode to configure an internet facing IP address as a VRRP critical-IP address, specifying the affected IP address, the associated VRID, and an optional decrement priority setting. A default ICMP probe is auto-configured to monitor remote critical-IP addresses. An administratively configured ICMP probe can be applied to override the default ICMP probe. See “[Preset Default ICMP Probes](#)” on page 12-6 for default ICMP probe details. Probes are configured in the tracked object manager. See [Chapter 12, Tracked Object Manager Configuration](#) for details.

If the critical-IP interface goes down with priority configured and enabled, the operational priority for the VRID to which this critical-IP address is associated is decremented by the value of the priority specified in this command. If the operational priority of the VRID falls below that of a backup router, the backup router becomes the master and the VRID assumes the priority value of the new master. Should the critical-IP interface come back up, the priority of the router associated with this critical-IP address is increased by the priority set for the critical-IP address. If preempt is enabled on the critical-IP address associated router, the router will once again become master and the VRID assumes the priority of the new master.


The default priority setting is enabled with a value of 10. Setting the critical-IP address priority to enabled signals that the critical-IP will affect the operational priority for the VRID. Setting the priority to disabled signals the critical-IP interface state will have no effect on the operational priority for the VRID.

Up to 2048 critical-IP addresses can be configured on a device. Up to 10 per VRID.

If the critical-IP address is configured on a router where the VRID IP address is owned by that router, the critical-IP configuration is ignored. When a router owns the IP address configured for the VRID, that router is automatically made the master with a hard-coded priority of 255. Only the failure of the interface with the VRID IP address can cause the router to move to backup status.

Figure 36-2 presents a typical critical-IP address configuration.

**Figure 36-2 Critical-IP Address Configuration**



The VRRP configuration is entered as follows:

- On both router 1 and router 2, in VLAN 111 configuration command mode, VRID 1 is created using the **vrrp create** command.

- On both router 1 and router 2, in VLAN 111 configuration command mode, the IP address 172.111.1.5 is assigned to VRID 1 using the **vrrp address** command.
- On router 1, in VLAN 111 configuration command mode, the VRRP priority is set to 105 using the **vrrp priority** command.
- On router 2, in VLAN 111 configuration command mode, the VRRP priority is set to 100 using the the **vrrp priority** command.
- On router 1, in VLAN 200 configuration command mode, configure IP address 172.200.1.1/16 as critical-IP address, enabling a priority of 10, using the **vrrp critical-ip** command.

In this example, should the critical-IP address 172.200.1.1/16 go down, the VRID 1 priority would decrement by 10, the value of the down critical-IP address, to 95. Router 2, with a priority set to 100 would take over as master. Should the critical-IP address 172.200.1.1/16 come back up, the priority for router 1 would increment by 10 from 95 to 105. Router 1 would now have a priority higher than the current priority 100 for VRID 1 and would become master once again.

## Configuring VRRP Authentication

A version 2 VRRP VRID can be configured for a simple clear text or encrypted MD5 authentication password. For MD5 authentication, 128-bit encryption is used unless the `hmac-96` option is specified, in which case 96-bit encryption is used.

Use the **vrrp authentication** command in interface configuration command mode, specifying the authentication type, a password, and in the case of MD5 optionally specifying 96-bit encryption.

## Enabling Master Preemption

By default, a router is enabled to preempt a lower priority master for the configured virtual router. If the router owns the virtual router IP address, it can not be preempted and always preempts other routers regardless of the priority setting or this preemption setting. Use the **vrrp preempt** command to enable or disable master preemption on this VRRP router.

It may be desirable to set a delay that a higher priority backup router should wait before preempting a lower priority master. By default there is no delay. To set a delay between 1 and 900 seconds, use the `vrrp preempt-delay` command in interface configuration command mode.

## Enabling Fabric Route Mode on the VRRP Backup Router



**Note:** Fabric route mode requires an advanced router license. See “[License Overview](#)” on page 9-10 for 7100-Series advanced router license details.


The fabric route mode allows a VRRP instance in the backup state to forward IPv4 and IPv6 packets destined for the VRRP MAC address. This feature provides for sharing of the traffic load across VRRP routers.

When fabric route mode is enabled on a VRRP router, the router behavior is exactly as detailed in RFC 5798 with these exceptions:

- **Accept-Mode** – A VRRP backup router with fabric route mode enabled will respond to ICMP packets it receives that are destined for the virtual gateway. A VRRP instance that wants to allow a non-owner virtual IP address to be pingable must enable the VRRP instance in fabric route mode to respond to ICMP packets using the **vrrp accept-mode** command.


- **Forwarding** – The VRRP standard forwarding defines a single VRRP master router and one or more VRRP backup router. Packets arriving on the backup router, destined for the VRRP MAC address, are cross bridged to the VRRP master router. The VRRP master router forwards the packet to the network (see [Figure 36-3](#) on page 6). The VRRP backup router remains in a dormant state and is not allowed to forward any traffic destined to the virtual gateway MAC address.

**Figure 36-3 Standard VRRP Forwarding**



With VRRP fabric route mode enabled on the backup router, a VRRP master and a backup router are started. Packets arriving on the backup router, destined to the VRRP MAC address, are no longer cross bridged. All IPv4 or IPv6 traffic destined for the VRRP MAC address and received directly on the VRRP backup router enabled for fabric route mode is now forwarded to the network by the backup router (See [Figure 36-4](#) on page 7). All ACL, policy rules, router application services, and routing is applied to the incoming packet.

Figure 36-4 Fabric Route VRRP Forwarding



- **IPV4 ARP** – All broadcast ARP requests are still handled by the VRRP master router except for unicast ARP requests destined for the VRRP MAC address received by the VRRP backup router with fabric route mode enabled. The VRRP backup fabric route mode enabled router processes these requests.
- **IPV6 Neighbor Discovery** – All neighbor solicitation requests are handled by the VRRP master router, except for unicast neighbor solicitation requests destined for the VRRP MAC address received by the VRRP fabric route mode enabled backup router. The VRRP fabric route mode enabled backup router will process these requests.

Use the **vrrp fabric-route-mode** command, in interface configuration mode, specifying the VRID the fabric mode is applied to, to enable fabric route mode on the VRRP router. This command can be applied to both the VRRP master and backup routers, but will only be operational on the VRRP backup router.

Use the **show running-config** command, specifying the interface the VRRP router is configured on, to display determine whether VRRP fabric route mode has been configured on the router.

## Enabling the VRRP Virtual Router

All other VRRP options must be set before enabling a VRRP virtual router on the routing interface. Once enabled, you can not make any configuration changes to VRRP without first disabling the interface, using the **no vrrp enable** command.

Use the **vrrp enable** command in interface configuration command mode, specifying the VRID of the virtual router to be enabled.

## Configuring VRRP

This section provides details for the configuration of VRRP on the 7100-Series products.



Table 36-1 lists VRRP parameters and their default values.

**Table 36-1 Default VRRP Parameters**

Parameter	Description	Default Value
accept-mode	Enables the master of this virtual router to accept IP packets for the configured IP address list, even if the device is not the owner.	disabled
advertise-interval	Specifies the interval between the advertisement the master sends to other routers participating in the selection process.	1 second
fabric route mode	A VRRP feature that allows a VRRP instance in the backup state to forward IPv4 and IPv6 packets destined for the VRRP MAC address.	disabled
interface-up delay	Specifies the delay in seconds for the VRID state transition from interface down to interface up	0 seconds (no delay)
priority	Specifies the router priority for the master election for this virtual router.	100
VRRP preemption	Specifies whether higher priority backup VRRP routers can preempt a lower priority master VRRP router and become master.	enabled

Procedure 36-1 describes how to configure VRRP. All VRRP configuration commands are entered in interface configuration command mode.

**Procedure 36-1 Configuring VRRP**

Step	Task	Command(s)
1.	Create a virtual router instance. Supported VRRP Versions: <ul style="list-style-type: none"> <li>• <b>v2-IPv4</b> - RFC 2338</li> <li>• <b>v3-IPv4</b> - draft-ietf-vrrp-unified-spec-03</li> <li>• <b>v3-IPv6</b> - draft-ietf-vrrp-unified-spec-03</li> </ul>	<b>vrrp create</b> <i>vrid version</i>
2.	Configure all VRRP IP addresses associated with this virtual router.	<b>vrrp address</b> <i>vrid ip-address</i> [ <b>enable</b>   <b>disable</b> ]
3.	Configure a VRRP primary address for this virtual router.	<b>vrrp primary-address</b> <i>vrid ip-address</i> [ <b>enable</b>   <b>disable</b> ]
4.	Optionally, change the VRRP router priority for this virtual router.	<b>vrrp priority</b> <i>vrid priority</i>
5.	Optionally, change the advertise interval for this virtual router.	<b>vrrp advertise-interval</b> <i>vrid</i> { <b>seconds</b> <i>interval</i>   <b>centiseconds</b> <i>interval</i> }
6.	Optionally, set the VRID state interface down to interface up delay.	<b>vrrp interface-up-delay</b> <i>vrid seconds</i>
7.	Configure any critical-IP interfaces for this virtual router.	<b>vrrp critical-ip</b> <i>vrid ip-address</i> [ <i>priority</i> ] [ <b>enable</b>   <b>disable</b> ] [ <b>remote</b> [ <i>probe-name</i> <i>probe-name</i> ]]

**Procedure 36-1 Configuring VRRP (continued)**

Step	Task	Command(s)
8.	Optionally, configure this router for VRRP authentication	<b>vrrp authentication</b> { <i>simple password</i>   <b>md5 password</b> [ <i>hmac-96</i> ]}
9.	Optionally, enable accept-mode for this virtual router, allowing the master to accept IP packets for the configured associated IP address list.	<b>vrrp accept-mode</b> <i>vrid</i>
10.	Optionally change the master preemption setting for this VRRP router.	<b>vrrp preempt</b> <i>vrid</i>
11.	Optionally, set the amount of time that will elapse before a backup VRRP router takes control from the current master when preemption is enabled.	<b>vrrp preempt-delay</b> <i>vrid delay</i>
12.	Optionally, in interface configuration mode, enable VRRP fabric route mode on the VRRP router.	<b>vrrp fabric-route-mode</b> <i>vrid</i>

Table 36-2 describes how to display VRRP information and statistics.

**Table 36-2 Displaying VRRP Information and Statistics**

Task	Command
Display VRRP configuration information for this system.	<b>show ip vrrp</b>
Display VRRP statistics for this system.	<b>show ip vrrp statistics</b>
Display VRRP configuration information for a specified interface.	<b>show ip vrrp interface</b> [ <i>vrid</i> ] [ <b>verbose</b> ]
Display detailed VRRP configuration information.	<b>show ip vrrp verbose</b>

## VRRP Configuration Examples

This section presents a two VRRP configuration examples:


- A basic VRRP configuration example with a single virtual router configured
- A multiple backup VRRP configuration with three virtual routers configured

### Basic VRRP Configuration Example

Figure 36-5 shows a basic VRRP configuration with a single virtual router. Routers R1 and R2 are both configured with one virtual router (VRID 1). Router R1 serves as the master because the VRRP router owns the IP address for this virtual router. Router R2 serves as the backup. The hosts are configured to use 172.111.1.1/16 as the default route.

The master advertise-interval is changed to 1.5 seconds for VRID 1.

If Router R1 should become unavailable, Router R2 would take over virtual router VRID 1 and its associated IP addresses. Packets sent to 172.111.1.1/16 would go to Router R2. When Router R1 comes up again, it would take over as master, and Router R2 would revert to backup.

**Figure 36-5 Basic Configuration Example****Router R1 Configuration of VRRP Instance 1**

```

7100 Chassis(rw)->configure
7100 Chassis(rw-config)->interface vlan 111
7100 Chassis(rw-config-intf-vlan.0.111)->ip address 172.111.1.1 255.255.255.0
primary
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp create 1 v3-IPv4
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp advertise-interval 1 centiseconds
150
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp address 1 172.111.1.1
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp enable 1
7100 Chassis(rw-config-intf-vlan.0.111)->no shutdown
7100 Chassis(rw-config-intf-vlan.0.111)->exit
7100 Chassis(rw-config)->show ip vrrp verbose
Interface: vlan.0.111
VRID: 1
  Version: 3, State: Master
  Master IP Address : 172.111.1.1
  Primary IP Address: 172.111.1.1
  Virtual MAC Address: 00:00:5E:00:01:01
  Advertisement Interval:  1.50 seconds
  Operational Priority:  255,  Configured Priority: 100
  Accept: no , Preempt: yes, Preempt time:  0 seconds
  Virtual IP Count:  1,  Critical IP Count:  0
  Virtual IP Addresses:
    172.111.1.1
  Critical IP Addresses:
    Interface                                Critical Priority      State
7100 Chassis(rw-config)->

```

**Router R2 Configuration of VRRP Instance 1**

```

7100 Chassis(rw)->configure
7100 Chassis(rw-config)->interface vlan 111

```

```

7100 Chassis(rw-config-intf-vlan.0.111)->ip address 172.111.1.2 255.255.255.0
primary
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp create 1 v3-IPv4
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp address 1 172.111.1.1
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp advertise-interval 1 centiseconds
150
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp enable 1
7100 Chassis(rw-config-intf-vlan.0.111)->no shutdown
7100 Chassis(rw-config-intf-vlan.0.111)->exit
7100 Chassis(rw-config)->show ip vrrp verbose
Interface: vlan.0.111
VRID: 1
  Version: 3, State: Backup
  Master IP Address : 172.111.1.1
  Primary IP Address: 172.111.1.2
  Virtual MAC Address: 00:00:6A:00:03:01
  Advertisement Interval:  1.50 seconds
  Operational Priority:  100, Configured Priority: 100
  Accept: no , Preempt: yes, Preempt time:  0 seconds
  Virtual IP Count:  1, Critical IP Count:  0
  Virtual IP Addresses:
    172.111.1.1
  Critical IP Addresses:
    Interface                                     Critical Priority      State
7100 Chassis(rw-config)->

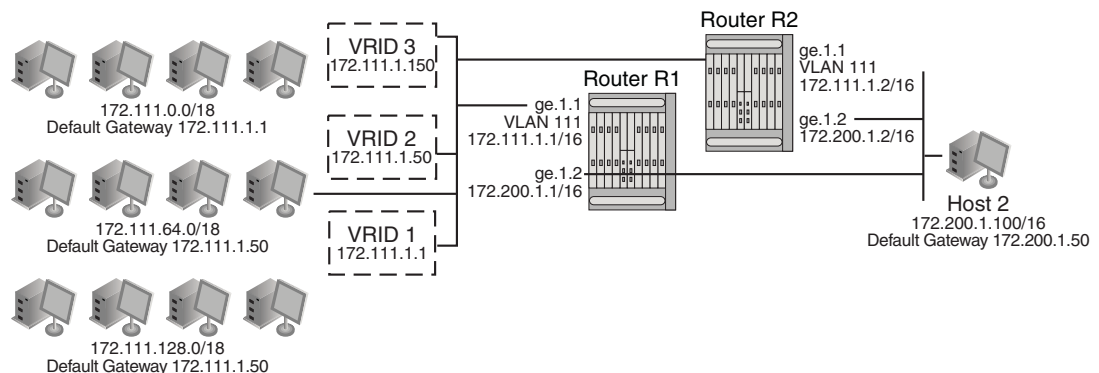
```

In this configuration, if an interface on VLAN 111 for Router R1 fails, the interface on Router R2 will take over for forwarding outside the local LAN segment.

## Multiple Backup VRRP Configuration Example

Figure 36-6 shows a multi-backup sample configuration.

**Figure 36-6 Multi-Backup VRRP Configuration Example**



Three VRRP instances are configured on VLAN 111 for both 7100-Series devices on Router R1's interface, 172.111.1.1, and Router R2's interface, 172.111.1.2. Each virtual router is given a different

virtual IP address that is used as a default gateway by a subset of hosts that reside of the LAN segment. Because interfaces on Router R1 and Router R2 for VLAN 111 are configured as belonging to VRID 1, 2, and 3, VRRP will support resiliency between these interfaces if one interface becomes in-operational.

To load balance traffic generated from the hosts on the 172.111.0.0/16 network, the hosts are partitioned into being configured with default gateways matching the virtual IP address of the VRRP virtual routers, and the VRRP Master for each VRRP instance is configured for distribution across Router R1 and Router R2. It is known that Router R1's interface, 172.111.1.1, will become Master for VRID 1 because it is the IP address owner for the virtual router. This interface is also configured to be Master for VRID 3 by raising its VRRP priority in VRRP instance 3 to 200. Therefore, Router R1's interface 172.111.1.1 will be Master for VRID 1 and VRID 3 handling traffic on this LAN segment sourced from subnets 172.111.0.0/18 and 172.111.128.0/18. Router R2's interface is configured to be the Master for VRID 2 by raising its VRRP priority in VRRP instance 2. Therefore, Router R2's interface 172.111.1.2 will be Master for VRID 2 handling traffic on this LAN segment sourced from subnets 172.111.64.0/18.

In this configuration, an interface on VLAN 111 for Router R1 or Router R2, or VRID 1, 2, or 3 fails, the interface on the other router will take over for forwarding outside the local LAN segment.

### Router R1

```
7100 Chassis(rw)->configure
7100 Chassis(rw-config)->interface vlan 111
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp create 1 v2-IPv4
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp address 1 172.111.1.1
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp enable 1
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp create 2 v2-IPv4
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp address 2 172.111.1.50
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp enable 2
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp create 3 v2-IPv4
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp address 3 172.111.1.150
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp priority 3 200
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp enable 3
7100 Chassis(rw-config-intf-vlan.0.111)->no shutdown
7100 Chassis(rw-config-intf-vlan.0.111)->exit
7100 Chassis(rw-config)->
```

### Router R2

```
7100 Chassis(rw)->configure
7100 Chassis(rw-config)->interface vlan 111
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp create 1 v2-IPv4
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp address 1 172.111.1.1
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp enable 1
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp create 2 v2-IPv4
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp address 2 172.111.1.50
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp priority 2 200
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp enable 2
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp create 3 v2-IPv4
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp address 3 172.111.1.150
7100 Chassis(rw-config-intf-vlan.0.111)->vrrp enable 3
```

```

7100 Chassis(rw-config-intf-vlan.0.111)->no shutdown
7100 Chassis(rw-config-intf-vlan.0.111)->exit
7100 Chassis(rw-config)->

```

## Terms and Definitions

[Table 36-3](#) lists terms and definitions used in this VRRP configuration discussion.

**Table 36-3 VRRP Configuration Terms and Definitions**

Term	Definition
Accept Mode	When enabled, it allows the master for this virtual router to accept IP packets for the configured associated IP address list.
Backup	The set of VRRP routers available to assume forwarding responsibility for a virtual router should the current Master fail.
IP Address owner	The VRRP router that has the virtual router's IP address(es) as real interface address(es). This is the router that, when up, will respond to packets addressed to one of these IP addresses for ICMP pings, TCP connections, etc.
Master	The VRRP router that is assuming the responsibility of forwarding packets sent to the IP address(es) associated with the virtual router, and answering ARP requests for these IP addresses.
Priority	<p>The priority field specifies the sending VRRP router's priority for the virtual router. Higher values equal higher priority. This field is an 8 bit unsigned integer field. The priority value for the VRRP router that owns the IP address(es) associated with the virtual router MUST be 255 (decimal).</p> <p>VRRP routers backing up a virtual router MUST use priority values between 1-254 (decimal). The default priority value for VRRP routers backing up a virtual router is 100 (decimal). The priority value zero (0) has special meaning indicating that the current Master has stopped participating in VRRP. This is used to trigger Backup routers to quickly transition to Master without having to wait for the current Master to timeout.</p>
Virtual Router	An abstract object managed by VRRP that acts as a default router for hosts on a shared LAN. It consists of a Virtual Router Identifier and a set of associated IP address(es) across a common LAN. A VRRP Router may backup one or more virtual routers.
VRID	Virtual Router ID — a unique number associated with each virtual router.
VRRP fabric route mode	A VRRP feature that allows a VRRP instance in the backup state to forward IPv4 and IPv6 packets destined for the VRRP MAC address.
VRRP Router	<p>A router running the Virtual Router Redundancy Protocol. It may participate in one or more virtual routers.</p> <p>A VRRP router may associate a virtual router with its real addresses on an interface, and may also be configured with additional virtual router mappings and priority for virtual routers it is willing to backup.</p>



## Security Configuration

This document provides the following information about configuring security features on the Extreme Networks 7100-Series platforms.

For information about...	Refer to page...
<a href="#">Using Security Features in Your Network</a>	<a href="#">37-1</a>
<a href="#">Implementing Security</a>	<a href="#">37-2</a>
<a href="#">Security Overview</a>	<a href="#">37-3</a>
<a href="#">Configuring Security</a>	<a href="#">37-8</a>

### Using Security Features in Your Network

The 7100-Series platform supports the following security features.

#### MAC Locking

The MAC locking security feature provides for limiting access to a port to specified MAC addresses or a maximum number of MAC addresses on a first come first serve basis. In the first case, a device with a MAC address that is not specifically configured will not be allowed access to a port. This provides the network administrator with confidence that only known devices will gain access to a port. The second case provides a means of controlling the maximum number of unique devices that will have access at any given time to the port configured for this mode of MAC locking.

#### Secure Shell

The Secure Shell (SSH) security feature provides a secure encrypted communications method between a client and the switch providing data privacy and integrity that is an alternative to the unsecure Telnet protocol. Using SSH the entire session is encrypted, including the transmission of user names and passwords, and negotiated between a client and server both configured with the SSH protocol. Telnet sessions are insecure. All data is sent unencrypted. Use SSH instead of Telnet when the security of login and data transmission is a concern.

The 7100-Series device supports both public key and password authentication methods.

#### TACACS+

The TACACS+ security feature provides an alternative to RADIUS for the authentication of devices desiring access to the network. TACACS+ provides device authentication, session



authorization, and per-command authorization, as well as accounting on a session and per command basis.

## Host Denial of Service (DoS)

The Host DoS security feature provides protection from all known attack vectors commonly used to deny service to the management entity of an Extreme Networks 7100-Series switch router. TCP, UDP and ICMP communications are monitored and reported on. Each attack type can be individually enabled and provides feed back in the form of display counters and SYSLOG messages when attacks are detected and prevented.

## Implementing Security

Take the following steps to implement supported 7100-Series security features in your network:

- To implement MAC locking:
  - Enable MAC locking both globally and on the ports to be configured for MAC locking
  - For ports that you are going to restrict access based upon a device's MAC address, set the port to MAC lock static and specify the maximum number of configure MAC addresses for that port
  - For ports you are going to restrict on a first come first serve for a set number of MAC addresses, enable dynamic MAC locking specifying the maximum number of MAC addresses allowed for that port
  - Optionally move all current dynamically enabled MAC locking MAC addresses to a static MAC locking configuration
  - Optionally allow dynamic MAC addresses to age based upon the configured MAC agetime
  - Optionally enable MAC lock trap messaging
- To implement Secure Shell:
  - Enable the SSH server
  - Set or reinitialize the host key
  - Verify the SSH state
- To implement TACACS+:
  - Enable TACACS+ on the client
  - Configure the TACACS+ server to be used by the client
  - Optionally enable TACACS+ session accounting
  - Optionally configure the TACACS+ session authorization service or privilege level
  - Optionally enable per command authorization
  - Optionally enable the TCP single connection feature for this device
- To implement Host DoS:
  - Enable one or more DoS attack mitigation types
  - Optionally set a logging event rate per a specified amount of time
  - Optionally enable logging
  - Verify the Host DoS configuration

## Security Overview

For information about...	Refer to page...
<a href="#">MAC Locking</a>	<a href="#">37-3</a>
<a href="#">Secure Shell</a>	<a href="#">37-4</a>
<a href="#">TACACS+</a>	<a href="#">37-5</a>
<a href="#">Host DoS</a>	<a href="#">37-7</a>

### MAC Locking

MAC Locking, sometimes referred to as MAC-based port locking, port locking, or port security, helps prevent unauthorized access to the network by limiting access based on a device's MAC address. MAC locking locks a port to one or more MAC addresses, preventing connection of unauthorized devices via a port. With MAC locking enabled, the only frames forwarded on a MAC locked port are those with the configured or dynamically selected MAC addresses for that port.


There are two different types of MAC locking:

- Static MAC Locking - Locking one or more specified MAC addresses to a port.
- Dynamic MAC Locking - Locking one or more MAC addresses to a port based on first arrival of received frames after dynamic MAC locking is enabled. The configuration specifies the maximum number of end users that will be allowed. As each new end user is identified, it is MAC locked up to the maximum number of users. Once the maximum number of users have been MAC locked, all other users will be denied access to the port until a MAC locked address is either aged, if aging is configured, or the session for that user ends.

MAC Locking is disabled by default. MAC locking must be both globally enabled and enabled on the desired ports. When globally enabling MAC lock you can optionally specify the port or ports to enable, or enable MAC locking on all ports. Once enabled, ports can be configured for either static or dynamic MAC locking. When configuring static MAC locking, specify the user MAC address and the port string for that user. When configuring dynamic MAC locking, specify the port and the maximum number of users that will be dynamically MAC locked. MAC addresses that are currently dynamically active can be auto reconfigured as static using the **set maclock move** command for the specified port.

Dynamic MAC lock address aging can be enabled per port. If the Filter DataBase (FDB) entry ages out for this station, the corresponding dynamic MAC locked stations will no longer be MAC locked. The age time for the FDB is set by the **set mac agetime** command and is displayed using the **show mac agetime** command. Dynamic MAC lock address aging is disabled by default.

[Figure 37-1](#) displays two users on a shared hub connected to an 7100-Series switch port. Data from the MAC locked user is forwarded on to the enterprise network. Data from the unconfigured user is dropped.

**Figure 37-1 Blocking Unauthorized Access with MAC Locking**

## Secure Shell

Secure Shell (SSH) is a protocol for secure remote login over an insecure network. SSH provides a secure substitute for Telnet by encrypting communications between two hosts.

The 7100-Series SSHv2 implementation includes:

- Data privacy
- Communication integrity

An SSH server resides on the 7100-Series platform and listens for client connection requests. Once a request is authenticated, a secure connection is formed through which all subsequent traffic is sent. All traffic is encrypted across the secure channel, which ensures data integrity. This prevents someone from seeing clear text passwords or file content, as is possible with the Telnet application.

Once SSH has been enabled and the 7100-Series has at least one valid IP address, you can establish an SSH session from any TCP/IP based node on the network, by using SSH to connect to an IP address, and entering your user name and password. Refer to the instructions included with your SSH application for information about establishing a session.

SSH is activated by enabling the SSH server on the device. Enabling the server automatically generates a host key for the server, used during the life of the client to server connection. The SSH server can be reinitialized. Reinitializing the server clears all current client to server connections. Reinitializing the server does not reinitialize the host key. Should you believe the host key has been compromised, or otherwise wish to change it, the host key can be reinitialized with a separate command.

During the handshake between an SSH client and an SSH server, each side sends a proposal of cryptographic Ciphers and Message Authentication Code (MAC)s. SSH Ciphers and MACs are applied to all new inbound (SSH server) and outbound (SSH client) SSH sessions. Existing sessions remain unchanged. Ciphers and MACs are entered in order of precedence from high to low.

Applied SSH Ciphers default to all supported ciphers in the following order of precedence: aes128-cbc, aes192-cbc, aes256-cbc, 3des-cbc, blowfish-cbc, and cast128-cbc.

When in FIPS mode, only the following FIPS compliant Ciphers are allowed (listed in the default order of precedence from high to low): aes128-cbc, aes192-cbc, aes256-cbc, and 3des-cbc. If non-FIPS Ciphers are configured when booting in FIPS mode, SSH uses the default Cipher list.

Use the **set ssh ciphers** command to administratively change the applied SSH Ciphers list. When using this command the order of precedence is modified to the order the Ciphers are entered. Any supported Cipher not entered is no longer allowed.

Applied MACs default to all supported MACs in the following order of precedence:  
hmac-sha1-etm @openssh.com, hmac-md5-etm @openssh.com,  
hmac-ripemd160-etm@openssh.com, hmac-sha1-96-etm @openssh.com, hmac-md5-96-etm @openssh.com, hmac-sha1, hmac-md5, hmac-ripemd160, hmac-ripemd160 @openssh.com, hmac-sha1-96, and hmac-md5-96.

When in FIPS mode, only the following FIPS compliant MACs are allowed (listed in the default order of precedence from high to low): hmac-sha1 and hmac-sha1-96. If non-FIPS MACs are configured when booting in FIPS mode, SSH uses the default MACs list.

Use the **set ssh macs** command to administratively change the applied SSH MACs list. When using this command the order of precedence is modified to the order the MACs are entered. Any supported MAC not entered is no longer allowed.

## TACACS+

TACACS+ (Terminal Access Controller Access Control System Plus) is a security protocol that can be used as an alternative to the standard RADIUS security protocol. The client function is implemented on the 7100-Series device to control access to this device in conjunction with a remote server. TACACS is defined in RFC 1492, and TACACS+ is defined in an un-published and expired Internet Draft draft-grant-tacacs-02.txt, "The TACACS+ Protocol Version 1.78", January, 1997.

TACACS+ client functionality falls into four basic capabilities: authentication and session authorization, per-command authorization, session accounting, and per-command accounting.

When the single connect feature is enabled, the TACACS+ client will use a single TCP connection for all requests to a given TACACS+ server.

## Session Authorization and Accounting

The TACACS+ client is disabled by default. When the TACACS+ client is enabled on the 7100-Series, using the **set tacacs enable** command, the session authorization parameters configured with the **set tacacs session authorization** command are sent by the client to the TACACS+ server when a session is initiated. The parameter values must match a service and access level attribute-value pairs configured on the server for the session to be authorized. If the parameter values do not match, the session will not be allowed. The service name and attribute-value pairs can be any character string, and are determined by your TACACS+ server configuration.

When session accounting is enabled, using the **set tacacs session accounting** command, the TACACS+ server will log accounting information, such as start and stop times, IP address of the remote user, and so forth, for each authorized client session. Once session accounting has been enabled, you can disable it with this command.

The 7100-Series device is informed of the TACACS+ server properties using the **set tacacs server** command. You can configure the timeout value for all configured servers or a single server, or you can configure the IP address, TCP port, and secret for a single server, specifying a server index value for this server.

## Per-Command Authorization and Accounting

In order for per-command accounting or authorization by a TACACS+ server to take place, the **set tacacs** command must be executed within an authorized session.

When per-command accounting is enabled, using the **set tacacs command accounting** command, the TACACS+ server will log accounting information, such as start and stop times, IP address of the client, and so forth, for each command executed during the session.

When per-command authorization is enabled, using the **set tacacs command authorization** command, the TACACS+ server will check whether each command is permitted for that authorized session and return a success or fail. If the authorization fails, the command is not executed.

## Single TCP Connection for All TACACS+ Requests

The 7100-Series device can be configured to use a single TCP connection for all TACACS+ client requests to a TACACS+ server. Use the **set tacacs singleconnect** command to enable this feature on the 7100-Series device.

## Host DoS

The Host DoS feature provides protection against all known DoS attack mitigation types.

Table 37-2 lists the configurable Host DoS mitigation types.

**Table 37-1 Host DoS Mitigation Types**

Threat	Description	Action
<b>Excessive Arp or ND</b>	Reception of an excessive number of ARP or ND frames from a single host.	Frames are discarded.
<b>Bad SIP</b>	Frames with a source IP address equal to multicast or broadcast.	Frames are discarded.
<b>Spoof</b>	Frames with a source IP address that is same as this router's interface address.	Frames are discarded.
<b>Christmas Tree</b>	Frames with an invalid TCP flag combination.	SYN+FIN and SYN+RST frames are discarded.
<b>Fragmented ICMP</b>	ICMP packets are fragmented.	All ICMP fragmented packets are discarded.
<b>ICMP Flood</b>	Excessive number of ICMP packets received.	Receipt of ICMP packets is limited to a user configurable limit of packets per second.
<b>Large ICMP</b>	ICMP packets exceed the configured maximum ICMP size.	ICMP packets exceeding the configured maximum ICMP size are discarded.
<b>Multicast/Broadcast Source address</b>	Packets with a Multicast or Broadcast source IP address.	Packets with a Multicast or Broadcast source IP address are discarded.
<b>LANd</b>	Packets with the destination IP address equal to the source IP address.	Packets with the destination IP address equal to the source IP address are discarded.
<b>Smurf</b>	A vulnerability due to ICMP directed broadcast packets.	ICMP directed broadcast packets are discarded.
<b>Fraggle Attack</b>	A vulnerability due to UDP directed broadcast packets.	UDP directed broadcast packets are discarded.
<b>SYN Flood</b>	Packets exceeding the maximum value and maximum establishment rate per source IP address or regardless of source.	Packets beyond established rates are discarded.
<b>Port Scan</b>	Packets exceeding the maximum value and maximum establishment rate.	Packets beyond established rates are discarded.

Globally enable host DoS for this device using the **hostdos enable** command. Host DoS is globally enabled by default. Entering a command line for each threat, specify the mitigation-type, in the **hostdos** command in global configuration command mode, to enable the specific DoS attack type to be mitigated.

The ICMP maximum allowed length can be set using the **hostdos** command **icmp-maxlength** option.

## Configuring Security

Table 37-2 lists Security parameters and their default values.

**Table 37-2 Default Security Parameters**

Parameter	Description	Default Value
MAC locking status	Specifies whether MAC locking is enabled or disabled both globally and on a specific port.	disabled
maximum number of dynamic MAC addresses	Specifies the maximum number of MAC addresses that will be locked on a port configured for dynamic MAC locking.	600
first arrival MAC address aging	Specifies that dynamic MAC locked addresses will be aged after the time set by the MAC agetime configuration.	disabled
MAC lock traps	Specifies whether traps associated with MAC locking will be sent.	disabled
maximum number of static MAC addresses	Specifies the maximum number of static MAC addresses allowed on a port.	64
SSH state	Specifies whether the SSH protocol is enabled or disabled on this device.	disabled
TACACS+ state	Specifies whether the TACACS+ protocol is enable or disabled on this device.	disabled
TACACS+ server timeout	Specifies the TACACS+ server timeout for the all TACACS+ servers.	10 seconds
session privilege level	Specifies the attribute value for the TACACS+ session management privilege level.	read-only = 0 read-write = 1 super-user = 15
TACACS+ single connect state	Specifies whether the TACACS+ single connect feature is enabled or disabled.	disabled

## Configuring MAC Locking

Procedure 37-1 describes how to configure MAC locking on an 7100-Series device. All MAC locking commands can be entered in any command mode.

**Procedure 37-1 MAC Locking Configuration**

Step	Task	Command(s)
1.	Globally enable MAC locking, optionally specifying the port(s) to be enabled. If no port is specified, all ports on the device are enabled. If one or more ports are specified, all unspecified ports remain disabled.	<b>set maclock enable</b> <i>[port_string]</i>

**Procedure 37-1 MAC Locking Configuration (continued)**

Step	Task	Command(s)
2.	Optionally, enable static MAC locking configuration on the specified port for the maximum number of MAC addresses specified by <i>value</i> .	<b>set maclock static</b> <i>port_string value</i>
3.	Optionally, create static MAC locking entries for the specified MAC address and port.	<b>set maclock mac_address port_string</b> {create   enable   disable}
4.	Optionally, create a dynamic MAC locking configuration, specifying the maximum number MAC addresses allowed for the specified port.	<b>set maclock firstarrival</b> <i>port_string value</i>
5.	Optionally, move all current dynamic MAC locking configured MACs to static entries.	<b>set maclock move</b> <i>port-string</i>
6.	Optionally, enable or disable first arrival MAC address aging on the specified port(s).	<b>set maclock agefirstarrival</b> <i>port_string</i> {enable   disable}
7.	Optionally, enable or disable MAC lock trap messaging.	<b>set maclock trap</b> <i>port_string</i> {enable   disable}

[Table 37-3](#) describes how to manage MAC locking on an 7100-Series port. All MAC locking commands can be entered in any command mode.

**Table 37-3 Managing MAC Locking**

Step	Task	Command(s)
1.	Display MAC locking configuration information for dynamic configurations, static configurations or by port.	<b>show maclock</b> [stations [firstarrival   static]] [ <i>port_string</i> ]
2.	Clear dynamic MAC locking configuration by port.	<b>clear maclock firstarrival</b> <i>port-string</i>
3.	Clear static MAC locking configuration by port.	<b>clear maclock static</b> <i>port_string</i>
4.	Clear MAC locking from one or more static MAC addresses.	<b>clear maclock</b> {all   <i>mac-address</i> } <i>port-string</i>

**MAC Locking Configuration Example**

The following command line enables MAC locking both globally for the device and at the port level for ports **ge.1.1** through **5**:

```
7100 Chassis(rw)->set maclock enable ge.1.1-5
7100 Chassis(rw)->
```

The following command lines enable **port ge.1.1** for a maximum of **3** static MAC address entries. This is followed by four static MAC address creation entries. The fourth entry fails because the maximum allowed has been set to 3:

```
7100 Chassis(rw)->set maclock static ge.1.1 3
7100 Chassis(rw)->set maclock 00-10-a4-e5-08-4e ge.1.1 create
7100 Chassis(rw)->set maclock 08-00-20-7c-e0-db ge.1.1 create
7100 Chassis(rw)->set maclock 00-60-08-14-4b-15 ge.1.1 create
7100 Chassis(rw)->set maclock 00-01-f4-2c-ad-b4 ge.1.1 create
Set failed for ge.1.1.
```



```
7100 Chassis(rw)->show maclock stations static
```

Port Number	MAC Address	Status	State	Aging
ge.1.1	00-10-a4-e5-08-4e	active	static	false
ge.1.1	00-60-08-14-4b-15	active	static	false
ge.1.1	08-00-20-7c-e0-db	active	static	false

```
7100 Chassis(rw)->
```

The following command lines configure ports **ge.1.2** through **5** for dynamic MAC locking with a maximum of **15** users on each port. This line is followed by a line enabling MAC locking trap messaging on ports **ge.1.1** through **5**:

```
7100 Chassis(rw)->set maclock firstarrival ge.1.2-5 15
```

```
7100 Chassis(rw)->set maclock trap ge.1.1-5 enable
```

```
7100 Chassis(rw)->
```

## Configuring Secure Shell

[Procedure 37-2](#) describes how to configure Secure Shell on an 7100-Series device. Secure Shell commands can be entered in any command mode.

### Procedure 37-2 SSH Configuration

Step	Task	Command(s)
1.	Enable, disable, or reinitialize the SSH server.	<b>set ssh {enable   disable   reinitialize}</b>
2.	Set or reinitialize the host key on the SSH server.	<b>set ssh hostkey [reinitialize]</b>
3.	Modify the SSH Ciphers list for all future sessions on this system.	<b>set ssh ciphers {aes128-cbc   aes192-cbc   aes256-cbc   3des-cbc   blowfish-cbc   cast128-cbc   rijndael-cbc@lysator.liu.se}</b>
4.	Modify the SSH MACs list for all future sessions on this system.	<b>set ssh macs {hmac-sha1-etm@openssh.com   hmac-md5-etm@openssh.com   hmac-ripemd160-etm@openssh.com   hmac-sha1-96-etm@openssh.com   hmac-md5-96-etm@openssh.com   hmac-sha1   hmac-md5   hmac-ripemd160   hmac-ripemd160@openssh.com   hmac-sha1-96   hmac-md5-96}</b>
5.	Verify the SSH state.	<b>show ssh state</b>

### SSH Configuration Example

The following commands enable and verify SSH:

```
7100 Chassis(rw)->set ssh enable
```

```
7100 Chassis(rw)->show ssh state
```

```
SSH Server state: Enabled
```

```
7100 Chassis(rw)->
```

The following command reinitializes the host key on the SSH server:

```
7100 Chassis(rw)->set ssh hostkey reinitialize
```

## Configuring TACACS+

Procedure 37-3 describes how to configure TACACS+ on an 7100-Series device. TACACS+ commands can be entered in any command mode.

### Procedure 37-3 TACACS+ Configuration

Step	Task	Command(s)
1.	Enable or disable the TACACS+ client.	<b>set tacacs {enable   disable}</b>
2.	Configure the TACACS+ server(s) to be used by the TACACS+ client.	<b>set tacacs server {index [ipaddress port [secret]]   all timeout timeout}</b>
3.	Optionally, enable TACACS+ session accounting	<b>set tacacs session accounting enable</b>
4.	Optionally, configure the TACACS+ session authorization service or privilege level. The attribute for privilege level is: <b>priv-lvl</b> .	<b>set tacacs session {authorization service name   read-only attribute value   read-write attribute value   super-user attribute value}</b>
5.	Optionally, enable per command accounting within an authorized session.	<b>set tacacs command accounting enable</b>
6.	Optionally, enable per command authorization.	<b>set tacacs command authorization enable</b>
7.	Optionally, enable the TCP single connection feature for this device.	<b>set tacacs singleconnect enable</b>

Table 37-4 describes how to manage TACACS+ on an 7100-Series device. All TACACS+ commands can be entered in any command mode.

### Table 37-4 Managing TACACS+

Task	Command(s)
Display TACACS+ configuration or state.	<b>show tacacs [state]</b>
Display the current TACACS+ server configuration.	<b>show tacacs server {index   all}</b>
Clear the TACACS+ server configuration or reset the server timeout to the default value.	<b>clear tacacs server {all   index} [timeout]</b>
Display the current TACACS+ client session settings.	<b>show tacacs session {authorization   accounting} [state]</b>
Reset TACACS+ session authorization settings to their default values.	<b>clear tacacs session authorization { [service] [read-only] [read-write] [super-user] }</b>
Display the current TACACS+ single connect state.	<b>show tacacs singleconnect [state]</b>

## TACACS+ Configuration Example

The following command enables TACACS+ on the TACACS+ client for this device:

```
7100 Chassis(rw)->set tacacs enable
```

The following commands configure and verify two TACACS servers for this device to indexes 1 and 2. Index 1 has an IP address of **10.10.10.20** on port **49** with a secret **mysecret1**. Index 2 has an IP address of **10.10.10.30** on port **49** with a secret of **mysecret2**. The server timeout value will remain at the default of 10 seconds.

```
7100 Chassis(rw)->set tacacs server 1 10.10.10.20 49 mysecret1
```

```
7100 Chassis(rw)->set tacacs server 2 10.10.10.30 49 mysecret2
```

```
7100 Chassis(rw)->show tacacs server all
```

```
TACACS+ Server IP Address Port Timeout Status
-----
1 10.10.10.20 49 10 Active
2 10.10.10.30 49 10 Active
```

7100 Chassis(rw)->

The following command enables and verifies session authorization for the exec service:

7100 Chassis(rw)->set tacacs session authorization service exec

7100 Chassis(rw)->show tacacs session authorization

```
TACACS+ service: exec
TACACS+ session authorization A-V pairs:
access level attribute value
read-only 'priv-lvl' '0'
read-write 'priv-lvl' '1'
super-user 'priv-lvl' '15'
```

7100 Chassis(rw)->

The following commands enable and verify session accounting, followed by commands that enable both accounting and authorization on a per command basis, for this device:

7100 Chassis(rw)->set tacacs session accounting enable

7100 Chassis(rw)->show tacacs session accounting

```
TACACS+ session accounting state: enabled
```

7100 Chassis(rw)->set tacacs command accounting enable

7100 Chassis(rw)->set tacacs command authorization enable

7100 Chassis(rw)->

The following command enables the TCP single connection feature for this device:

7100 Chassis(rw)->set tacacs singleconnect

7100 Chassis(rw)->

## Configuring Host DoS

[Procedure 37-4](#) describes how to configure Host DoS on an 7100-Series device. Host DoS configuration commands are entered in global configuration command mode.

### Procedure 37-4 Host DoS Configuration

Step	Task	Command(s)
1.	Enable host DoS globally for this device. Threats must be specifically enabled for mitigation to occur for that threat as specified in the following step.	<b>hostdos enable</b>
2.	Enable a mitigation type, and optionally set the rate at which events will be acted upon.	<b>hostdos {mitigation-type   enable   icmp-maxlength icmp-maxlength} [rate count [per-second   per-minute   per-hour   per-day]] [nolog]</b>
3.	Optionally, disable logging for the specified DoS attack types.	<b>hostdos mitigation-type nolog</b>
4.	Optionally, specify an ICMP maximum packet size for <b>icmpsize</b> mitigation.	<b>hostdos icmpsize maxlength length</b>

Table 37-4 describes how to display Host DoS configuration state and counters on an 7100-Series device.

**Table 37-5 Displaying Host DoS**

Step	Task	Command(s)
1.	Display configuration state for one or all Host DoS attack mitigation types.	<b>show hostdos</b> <i>[mitigation-type]</i>
2.	Display statistic counters for one or all Host DoS attack mitigation types.	<b>show hostdos</b> <i>[mitigation-type]</i> <b>[stats]</b>

## Host DoS Configuration Example

This example shows how to:

- Globally enables host Dos on this device
- Enable the checkSpoof mitigation type, with a log display rate of 5 per-minute
- Enable the XmasTree mitigation type and disable logging for this threat

```
7100 Chassis(rw-config)->hostdos enable
7100 Chassis(rw-config)->hostDoS spoof rate 5 per-minute
7100 Chassis(rw-config)->hostdos xmastree nolog
7100 Chassis(rw-config)->show hostDoS
hostDoS is globally enabled
hostDoS icmp-maxlength is 1024
hostDoS Spoof is enabled , logging is enabled , rate is 5 per-minute
hostDoS XmasTree is enabled , logging is disabled, rate is 0 per-second
hostDoS IcmpFrag is disabled, logging is enabled , rate is 0 per-second
hostDoS IcmpFlood is disabled, logging is enabled , rate is 0 per-second
hostDoS IcmpSize is disabled, logging is enabled , rate is 0 per-second
hostDoS BadSIP is disabled, logging is enabled , rate is 0 per-second
hostDoS LANd is disabled, logging is enabled , rate is 0 per-second
hostDoS Smurf is disabled, logging is enabled , rate is 0 per-second
hostDoS Fraggle is disabled, logging is enabled , rate is 0 per-second
hostDoS SynFlood is disabled, logging is enabled , rate is 0 per-second
hostDoS PortScan is disabled, logging is enabled , rate is 0 per-second
hostDoS TearDrop is disabled, logging is enabled , rate is 0 per-second
7100 Chassis(rw-config)->
```



## Access Control List Configuration

This document provides the following information about configuring IPv4 and IPv6 Access Control Lists (ACLs) on the Extreme Networks 7100-Series platforms.

For information about...	Refer to page...
<a href="#">Using Access Control Lists (ACLs) in Your Network</a>	38-1
<a href="#">Implementing ACLs</a>	38-1
<a href="#">ACL Overview</a>	38-2
<a href="#">Configuring ACLs</a>	38-8
<a href="#">Terms and Definitions</a>	38-14

### Using Access Control Lists (ACLs) in Your Network

ACLs allow the configuration of permit and denial of IPv4 and IPv6 packet forwarding based upon IP address, protocol, and port matching, depending upon the ACL type. The 7100-Series firmware supports configuration of both standard and extended ACLs. Standard ACLs allow the packet source IP address to be configured, while extended ACLs allow both source and destination IP addresses, protocol and TCP or UDP port matching, as well as the optional specifying of a DSCP, ToS, or IP precedence value. ACLs are also used to match addresses or traffic by client applications such as route map (for policy-based routing and route redistribution) and IP Directed Broadcast.

### Implementing ACLs

To implement an ACL on your network:

- Create the ACL
- Enter the rules and comments for this ACL that will determine which packets will be forwarded or not forwarded on the routing interface this ACL will be applied to
- Optionally manage your ACL by:
  - Copying a preexisting ACL to a non-existing ACL
  - Appending a preexisting ACL to another preexisting ACL
  - Entering an ACL comment entry
  - Deleting an ACL rule entry
  - Inserting a new ACL rule entry into an ACL
  - Moving an ACL rule to a new location in an ACL

- Apply the ACL to a routing interface

## ACL Overview

This section describes ACL creation, rule entry, and application of the ACL to a routing interface required to implement an ACL, as well as, the features available for managing ACL rules and displaying ACLs.



**Note:** An "implicit deny" is hard coded at the end of all ACLs. The implicit deny blocks anything not explicitly permitted within the ACL, including routing protocols and management connections.

## Creating an ACL

There are two types of ACLs: standard and extended. The type of ACL you need depends exclusively upon the packet field(s) that will generate a hit for the rules specified in the ACL. For a standard ACL, only the source IP address is configurable. For an extended ACL, the protocol, source IP address, destination IP address, and in the case of the TCP or UDP protocols, matching source and destination ports are configurable.

There are two ways to identify the new ACL: a number or a name. The use of a number is for IPv4 ACLs only. Standard IPv4 ACL numbers range from **1** to **99**. Extended IPv4 ACL numbers range from **100** to **199**. Both IPv4 and IPv6 allow alphanumeric names that must start with an alpha character. A name may be quoted, as the quotes are stripped, but spaces are not supported in the quoted string. A name cannot be one of the **show access-lists** keywords **brief** or **applied**, or any prefix thereof such as **?br?** or **?app?**. Names can be up to 64 characters in length.

Once you have determined the appropriate ACL type, use the:

- **ip access-list standard** command to create an IPv4 standard access-list and **ipv6 access-list standard** command to create an IPv6 standard access-list
- **ip access-list extended** command to create an IPv4 extended access-list and **ipv6 access-list extended** command to create an IPv6 extended access-list

In each case, specifying the access-list number or name for the ACL.

An existing ACL can be copied to a non-existing ACL of the same IP type (IPv4 or IPv6). An existing ACL can be appended to the end of another existing ACL of the same IP type, but a standard ACL may not be appended to an extended ACL nor vice versa.

Upon creating the ACL, you are placed in the access-list configuration command mode where you can enter rules or comment entries for this ACL.

### IPv4 ACL Creation Examples

The following example creates a standard IPv4 ACL with the access-list number **1** as its identifier:

```
7100 Chassis(rw-config)->ip access-list standard 1
7100 Chassis(rw-cfg-std-acl)->
```

The following example creates an extended IPv4 ACL with the access-list number **100** as its identifier:

```
7100 Chassis(rw-config)->ip access-list extended 100
7100 Chassis(rw-cfg-ext-acl)->
```

The following example creates a standard ACL with the name **ipv4acl1** as its identifier:

```
7100 Chassis(rw-config)->ip access-list standard ipv4acl1
7100 Chassis(rw-cfg-std-acl)->
```

## IPv6 ACL Creation Examples

The following example creates a standard IPv6 ACL with the access-list number **acl1** as its identifier:

```
7100 Chassis(rw-config)->ipv6 access-list standard acl1
7100 Chassis(rw-cfg-ipv6-std-acl)->
```

The following example creates an extended IPv6 ACL with the access-list number **acl100** as its identifier:

```
7100 Chassis(rw-config)->ipv6 access-list extended 100
7100 Chassis(rw-cfg-ipv6-ext-acl)->
```

The following example creates a standard IPv6 ACL with the name **ipv6acl1** as its identifier:

```
7100 Chassis(rw-config)->ipv6 access-list standard ipv6acl1
7100 Chassis(rw-cfg-ipv6-std-acl)->
```

## Creating ACL Rules

ACL rules define the basis upon which a hit will take place for the ACL. Rules in an ACL are order-dependent. A packet is either forwarded (a **permit** rule) or not forwarded (a **deny** rule) according to the first rule that is matched. The matching criteria available is determined based upon whether the ACL is a standard ACL or an extended ACL. As soon as a rule is matched, processing of the access list stops. There is an implicit “deny all” rule at the end of every ACL. If all rules are missed, the packet is not forwarded.



**Note:** When applying an ACL to a non-host context, the following restrictions apply:

- Ranges are not allowed for UDP and TCP rules; the equal (eq) option is supported
- The **msg** option is not supported for ICMP rules



**Note:** The resources available for the configuration of ACL rules is based upon the system resource allocation profile. See “[System Resource Allocation Profile](#)” on page 22-9 for system resource allocation profile configuration details. See the release notes that come with your firmware for ACL resource limit details.

For an extended ACL, the following protocols can be specified in a rule:

- A specific or all internet protocols
- Authentication Header protocol
- Encapsulation Security Payload
- Generic Router Encapsulation protocol
- An established TCP connection
- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)
- Internet Control Message Protocol (ICMP or ICMPv6)

TCP and UDP rules can match source and destination ports against the following values: equal to. TCP rules can also distinguish established connections for new connection requests.

Extended ACLs can optionally be set for a Diffserv codepoint (DSCP), IP precedence, or IP Type of Service (ToS) value for both IPv4 and IPv6. IPv6 provides additional support for routing header match against source-routed packet, and the packet’s routing extension header, mobility extension header, and mobility-type extension header.



For a standard ACL, a source IPv4 address and an optional wildcard or IPv6 address and length are specified for the rule. For an extended ACL a source and destination IP address and wildcard are specified for the rule. In the case of an IPv4, Source and destination wildcards provide an inverted mask (specifies the don't care bits as 1s). 0.0.0.0 specifies an exact match. An **any** option is available for. The any option is short hand for 0.0.0.0 255.255.255.255.

Logging of ACL configuration activity is supported via syslog messages. This logging can be enabled for a specified entry, all entries, or the final implicit deny rule using the **log** entry command in access list configuration mode. Logging format can be in either a verbose or summary format.

Comments can be entered at the next available entry location, and, once entered, can be moved to a desired location.

Use the **permit** command to create a rule that forwards packets based upon the defined rule.

Use the **deny** command to create a rule that prevents the forwarding of packets based upon the defined rule.

### IPv4 ACL examples

The following example creates a standard ACL **1**, and specifies an entry 1 permit rule with a source IP address of 10.0.0.1 and a wild card of 0.0.255.255. The explicit deny all rule denies all other traffic for this ACL:

```
7100 Chassis(rw-config)->ip access-list standard 1
7100 Chassis(rw-cfg-std-acl)->permit 10.0.0.1 0.0.255.255
7100 Chassis(rw-cfg-std-acl)->show access-lists 1
Standard IP access list 1 (2 entries)
  1 permit 10.0.0.1 0.0.255.255
  -- implicit deny all --
```

The following example creates an extended access-list 120 and configures a deny entry for the IP protocol with a source address 20.0.0.1 and source wildcard of 0.0.255.255 and a destination address of any. Syslog messaging is enabled to log any hit for this rule. This rule is followed by a permit rule for any other source or destination IP protocol traffic:

```
7100 Chassis(rw-config)->ip access-list extended 120
7100 Chassis(rw-cfg-ext-acl)->deny ip 20.0.0.1 0.0.255.255 any log
7100 Chassis(rw-cfg-ext-acl)->permit ip any any
7100 Chassis(rw-cfg-ext-acl)->show access-lists 120
Extended IP access list 120 (3 entries)
  1 deny ip 20.0.0.1 0.0.255.255 any
  2 permit ip any any
  -- implicit deny all --
7100 Chassis(rw-cfg-ext-acl)->
```

### IPv6 ACL Examples

This example enters configuration mode for standard IPv6 access list **acl2** and configures a permit entry for source address 2001:1234:50:0:21f:45ff:fe3d:21be/64:

```
7100 Chassis(rw-config)->ipv6 access-list standard acl2
7100 Chassis(rw-cfg-ipv6-ext-acl)->permit 2001:1234:50:0:21f:45ff:fe3d:21be/64
7100 Chassis(rw-cfg-ipv6-ext-acl)->
```

This example enters configuration mode for extended IPv6 access list **acl120** and configures a permit entry for the IP protocol with a source address **2001:1234:50:0:21f:45ff:fe3d:21aa/64** and a destination address of any:

```
7100 Chassis(rw-config)->ipv6 access-list extended acl120
7100 Chassis(rw-cfg-ipv6-ext-acl)->permit ipv6
2001:1234:50:0:21f:45ff:fe3d:21aa/64 any
7100 Chassis(rw-cfg-ipv6-ext-acl)->
```

## Managing ACL Rules

Existing ACL rules can be deleted, moved, or replaced. New rules can be inserted at a specified location, otherwise rules are placed at the next available entry value. Comments can be entered into an ACL to provide useful information about the ACL. The contents of one or all ACLs can be displayed.

### Deleting an ACL Rule

An ACL rule or range of rules can be deleted using the **delete** command.

The following example displays an extended ACL 120 and deletes and deletes entries 2 and 3:

```
7100 Chassis(rw-config)->ip access-list extended 120
7100 Chassis(rw-cfg-ext-acl)->show access-lists 120
Extended IP access list 120 (5 entries)
  1 deny ip 20.0.0.1 0.0.255.255 any
  2 deny ip 30.0.0.1 0.0.255.255 any
  3 deny ip 40.0.0.1 0.0.255.255 any
  4 permit ip any any
  -- implicit deny all --
7100 Chassis(rw-cfg-ext-acl)->delete from 2 to 3
7100 Chassis(rw-cfg-ext-acl)->show access-lists 120
Extended IP access list 120 (3 entries)
  1 deny ip 20.0.0.1 0.0.255.255 any
  2 permit ip any any
  -- implicit deny all --
```

The following example enters configuration mode for standard IPv6 access list **acl2** and deletes rule entry **10 - 12**:

```
7100 Chassis(rw-config)->ipv6 access-list standard acl2
7100 Chassis(rw-cfg-ipv6-std-acl)->delete from 10 to 12
7100 Chassis(rw-cfg-ipv6-std-acl)->
```

### Moving an ACL Rule

An ACL rule or range of rules can be moved to a different location in the ACL using the **move before** command.

The following example displays an extended ACL **121** and moves entries **3** and **4** to before entry **2**:

```
7100 Chassis(rw-config)->ip access-list extended 121
7100 Chassis(rw-cfg-ext-acl)->show access-lists 121
Extended IP access list 121 (5 entries)
```

```
1 deny ip 20.0.0.1 0.0.255.255 any
2 permit ip any any
3 deny ip 30.0.0.1 0.0.255.255 any
4 deny ip 40.0.0.1 0.0.255.255 any
-- implicit deny all --
7100 Chassis(rw-cfg-ext-acl)->move before 2 from 3 to 4
7100 Chassis(rw-cfg-ext-acl)->show access-lists 121
Extended IP access list 121 (5 entries)
1 deny ip 20.0.0.1 0.0.255.255 any
2 deny ip 30.0.0.1 0.0.255.255 any
3 deny ip 40.0.0.1 0.0.255.255 any
4 permit ip any any
-- implicit deny all --
```

This example enters configuration mode for standard IPv6 access list **acl2** and moves rule entries **10 - 12** before rule entry **5**:

```
7100 Chassis(rw-config)->ipv6 access-list standard acl2
7100 Chassis(rw-cfg-ipv6-std-acl)->move before 5 from 10 to 12
7100 Chassis(rw-cfg-ipv6-std-acl)->
```

## Replacing an ACL Rule

An ACL rule or range of rules can be replaced by a specified permit, deny, or remark using the **replace** command.

The following example displays an extended ACL 121 and replaces entry 1 with a deny rule for source IP address 10.0.0.1 and destination IP address any:

```
7100 Chassis(rw-config)->ip access-list extended 121
7100 Chassis(rw-cfg-ext-acl)->show access-lists 121
Extended IP access list 121 (5 entries)
1 deny ip 20.0.0.1 0.0.255.255 any
2 deny ip 30.0.0.1 0.0.255.255 any
3 deny ip 40.0.0.1 0.0.255.255 any
4 permit ip any any
-- implicit deny all --
7100 Chassis(rw-cfg-ext-acl)->replace 1 deny ip 10.0.0.1 0.0.255.255 any
7100 Chassis(rw-cfg-ext-acl)->show access-lists 121
Extended IP access list 121 (5 entries)
1 deny ip 10.0.0.1 0.0.255.255 any
2 deny ip 30.0.0.1 0.0.255.255 any
3 deny ip 40.0.0.1 0.0.255.255 any
4 permit ip any any
-- implicit deny all --
```

This example replaces entry **1** of IPv6 access list **acl10** with a permit any source address :

```
7100 Chassis(rw-config)->ipv6 access-list standard acl10
7100 Chassis(rw-cfg-ipv6-std-acl)->replace 1 permit any
7100 Chassis(rw-cfg-ipv6-std-acl)->
```

## Inserting an ACL Rule

When entering an ACL rule, the new rule is appended to the end of the ACL by default. A new ACL rule can be inserted into a specified entry location using the **insert before** command.

The following example displays an extended ACL 121 and inserts a new entry 2 with a deny rule for source IP address 20.0.0.1 and destination IP address any:

```
7100 Chassis(rw-config)->ip access-list extended 121
7100 Chassis(rw-cfg-ext-acl)->show access-lists 121
Extended IP access list 121 (5 entries)
  1 deny ip 10.0.0.1 0.0.255.255 any
  2 deny ip 30.0.0.1 0.0.255.255 any
  3 deny ip 40.0.0.1 0.0.255.255 any
  4 permit ip any any
  -- implicit deny all --
7100 Chassis(rw-cfg-ext-acl)->insert before 2 deny ip 20.0.0.1 0.0.255.255 any
7100 Chassis(rw-cfg-ext-acl)->show access-lists 121
Extended IP access list 121 (6 entries)
  1 deny ip 10.0.0.1 0.0.255.255 any
  2 deny ip 20.0.0.1 0.0.255.255 any
  3 deny ip 30.0.0.1 0.0.255.255 any
  4 deny ip 40.0.0.1 0.0.255.255 any
  5 permit ip any any
  -- implicit deny all --
```

This example enters configuration mode for extended IPv6 access list **acl10** and inserts a rule before entry **10** that permits packets with a source address for host **2002:100::50** and a destination address of **2001:100::100:25/64** with a ToS value of **6**:

```
7100 Chassis(rw-config)->ipv6 access-list standard acl10
7100 Chassis(rw-cfg-ipv6-ext-acl)->insert before 10 permit host 2002:100::50
2001:100::100:25/64 traffic-class 6
7100 Chassis(rw-cfg-ipv6-ext-acl)->
```

## Applying ACLs

Once you have defined an ACL, it can be applied per routing interface. An ACL can be applied to host access or an interface before it is created. The association of the name or number of the ACL to the host or interface is persistent. You can use ACLs to filter traffic on individual interfaces, with a directional context (inbound, outbound, or both).

Use the **ip access-group** command to apply an IPv4 access-list to an interface and the **ipv6 access-group** command to apply an IPv6 access-list to an interface, in interface configuration command mode, specifying the access-list number or name followed by the directional context to which this ACL will be applied.

Use the **ip host-access** command for an IPv4 access-list and the **ipv6 host-access** command for an IPv6 access-list in configuration command mode, specifying the access-list number or name, to apply an ACL to host services for this device.

Use the **show access-lists applied** to display access-lists that have been applied to a routing interface.

The following example applies the extended ACL 121 to both the inbound and outbound direction on VLAN 2.

```
7100 Chassis(su-config)->interface vlan 2
7100 Chassis(su-config-intf-vlan.0.2)->ip access-group 121 in
7100 Chassis(su-config-intf-vlan.0.2)->ip access-group 121 out
7100 Chassis(su-config-intf-vlan.0.2)->show access-lists applied
Extended IP access list 121, applied inbound on interface 2 (5 entries)
Extended IP access list 121, applied outbound on interface 2 (5 entries)
7100 Chassis(su-config-intf-vlan.0.2)->
```

This example shows how to apply the standard access list `acl10` for all inbound frames on VLAN 50. Based upon the definition of access list `acl10`, only frames with source `fe80:0:0:0:21f:45ff:fe3d:21aa/64` are routed. All the frames with other sources received on VLAN 50 are dropped:

```
7100 Chassis(su-config)->ipv6 access-list standard acl10
7100 Chassis(su-cfg-ipv6-std-acl)->permit fe80:0:0:0:21f:45ff:fe3d:21aa/64 log
7100 Chassis(su-cfg-ipv6-std-acl)->exit
7100 Chassis(su-config)->interface vlan 50
7100 Chassis(su-config-intf-vlan.0.50)->ipv6 access-group acl10 in
7100 Chassis(su-config-intf-vlan.0.50)->
```

## Permit Layer 4 Fragmented Packet Behavior

In filtering traffic a packet is tested against each rule in the ACL. A rule in an extended ACL may specify layer four information such as TCP/UDP ports or ICMP type, code, or TCP flags for an established connection. If packets are sent as a series of fragments, Layer 4 information is present in the initial fragment, but is not present in subsequent fragments. Although a subsequent fragment should match the same rule matched by the initial fragment, all the information is not available at the time the rule is tested.

To deal with this problem for subsequent fragments, permit rules are handled differently than deny rules. The philosophy is that is better to leak fragments that should be dropped than to drop fragments that should be forwarded, as the latter would break a receiving application.

Consider a subsequent fragment that matches all the values specified in a rule, except for the Layer 4 information. If it is a permit rule it is considered to match the rule; that fragment is forwarded and no further ACL rules are checked. If a deny rule, the fragment is considered to not match and its fate depends on later rules. For example, it would be dropped if it got all the way to the implicit deny or it would be forwarded if it matched a later permit rule that specified only the same layer three information.

## Configuring ACLs

This section provides details for the configuration of ACLs on the 7100-Series products.

[Procedure 38-1](#) describes how to create an IPv4 ACL and manage IPv4 ACLs at the ACL level.

**Procedure 38-1 Creating and Managing IPv4 and IPv6 ACLs**

Step	Task	Command(s)
1.	In global configuration command mode, create a standard or extended IPv4 or IPv6 ACL, or enter IPv4 or IPv6 ACL configuration mode for an already existing ACL.	<b>ipv4 access-list</b> { <b>standard</b>   <b>extended</b> } { <i>access-list-number</i>   <i>name</i> }  <b>ipv6 access-list</b> { <b>standard</b>   <b>extended</b> } <i>name</i>
2.	In global configuration command mode, optionally, copy a preexisting IPv4 or IPv6 ACL to a non-existing IPv4 or IPv6 ACL.	<b>ipv4 ip access-list</b> { <b>standard</b>   <b>extended</b> } { <i>access-list-number</i>   <i>name</i> } <b>copy to</b> { <i>access-list-number</i>   <i>name</i> }  <b>ipv6 ip access-list</b> { <b>standard</b>   <b>extended</b> } <i>name</i> <b>copy to</b> <i>name</i>
3.	In global configuration command mode, optionally, append a preexisting IPv4 or IPv6 ACL to another preexisting IPv4 or IPv6 ACL.	<b>ipv4 ip access-list</b> { <b>standard</b>   <b>extended</b> } { <i>access-list-number</i>   <i>name</i> } <b>append to</b> { <i>access-list-number</i>   <i>name</i> }  <b>ipv6 ip access-list</b> { <b>standard</b>   <b>extended</b> } <i>name</i> <b>append to</b> <i>name</i>
4.	In global configuration command mode, optionally, check the efficiency of an IPv4 or IPv6 ACL.	<b>ipv4 ip access-list</b> { <b>standard</b>   <b>extended</b> } { <i>access-list-number</i>   <i>name</i> } <b>check</b>  <b>ipv6 ip access-list</b> { <b>standard</b>   <b>extended</b> } <i>name</i> <b>check</b>

[Procedure 38-2](#) describes how to enter and manage standard ACL rules.

**Procedure 38-2 Entering and Managing Standard IPv4 ACL Rules**

Step	Task	Command(s)
1.	In IPv4 ACL configuration command mode, optionally, create a standard IPv4 ACL deny rule entry.	<b>deny</b> { <i>source source-wildcard</i>   <b>any</b>   <b>host</b> <i>ip-address</i> } [ <b>log</b>   <b>log-verbose</b> ]
2.	In IPv4 ACL configuration command mode, optionally, insert a new standard IPv4 ACL rule entry before the specified preexisting entry for this standard ACL.	<b>insert before</b> <i>entry</i> { <b>remark</b> " <i>text</i> "   { <b>permit</b>   <b>deny</b> } { <i>source source-wildcard</i>   <b>any</b>   <b>host</b> <i>ip-address</i> } [ <b>log</b>   <b>log-verbose</b> ]}
3.	In IPv4 ACL configuration command mode, optionally, replace the specified standard ACL entry with the specified new entry.	<b>replace</b> <i>entry</i> { <b>remark</b> " <i>text</i> "   <b>deny</b> { <i>source</i> [ <i>source-wildcard</i> ]   <b>any</b>   <b>host</b> <i>ip-address</i> ]   <b>permit</b> { <i>source</i> [ <i>source-wildcard</i> ]   <b>any</b>   <b>host</b> <i>ip-address</i> }}

[Procedure 38-3](#) describes how to enter and manage standard ACL rules.

**Procedure 38-3 Entering and Managing Standard IPv6 ACL Rules**

Step	Task	Command(s)
1.	In IPv6 ACL configuration command mode, optionally, create a standard IPv6 ACL permit rule entry.	<b>permit</b> { <i>source-address/length</i>   <b>any</b>   <b>host</b> <i>ip-address</i> } [ <b>log</b>   <b>log-verbose</b> ]
2.	In IPv6 ACL configuration command mode, optionally, create a standard IPv6 ACL deny rule entry.	<b>deny</b> { <i>source-address/length</i>   <b>any</b>   <b>host</b> <i>ip-address</i> } [ <b>log</b>   <b>log-verbose</b> ]

**Procedure 38-3 Entering and Managing Standard IPv6 ACL Rules (continued)**

Step	Task	Command(s)
3.	In IPv6 ACL configuration command mode, optionally, insert a new standard IPv6 ACL rule entry before the specified preexisting entry for this standard ACL.	<b>insert before entry</b> { <b>remark text</b>   { <b>permit</b>   <b>deny</b> }} { <i>source-address/length</i>   <b>any</b>   <b>host ip-address</b> } [ <b>log</b>   <b>log-verbose</b> ]
4.	In IPv6 ACL configuration command mode, optionally, replace the specified standard ACL entry with the specified new entry.	<b>replace entry</b> { <b>remark text</b>   { <b>permit</b>   <b>deny</b> }} { <i>source-address/length</i>   <b>any</b>   <b>host ip-address</b> } [ <b>log</b>   <b>log-verbose</b> ]

[Procedure 38-4](#) describes how to enter and manage extended IPv4 ACL rules.

**Procedure 38-4 Entering and Managing Extended IPv4 ACL Rules**

Step	Task	Command(s)
1.	In IPv4 ACL configuration command mode, optionally, create an extended IPv4 ACL permit rule entry.	<p><b>permit</b> {<i>protocol-num</i>   <b>ip</b>   <b>ah</b>   <b>esp</b>   <b>gre</b>} {<i>source source-wildcard</i>   <b>any</b>   <b>host ip-address</b>} {<i>destination destination-host wildcard</i>   <b>any</b>   <b>host ip-address</b>} [<b>dscp code</b>] [<b>precedence value</b>] [<b>tos value</b>] [<b>log</b>   <b>log-verbose</b>]</p> <p><b>permit tcp</b> {<i>source source-wildcard</i>   <b>any</b>   <b>host ip-address</b>} [{<b>eq</b>   <b>neq</b>   <b>gt</b>   <b>lt</b>} <i>source-port</i>] [<b>range start-port end-port</b>] {<i>destination destination-host wildcard</i>   <b>any</b>   <b>host ip-address</b>} [{<b>eq</b>   <b>neq</b>   <b>gt</b>   <b>lt</b>} <i>dest-port</i>] [<b>range start-port end-port</b>] [<b>established</b>] [<b>dscp code</b>] [<b>precedence value</b>] [<b>tos value</b>] [<b>log</b>   <b>log-verbose</b>]</p> <p><b>permit udp</b> {<i>source source-wildcard</i>   <b>any</b>   <b>host ip-address</b>} [{<b>eq</b>   <b>neq</b>   <b>gt</b>   <b>lt</b>} <i>source-port</i>] [<b>range start-port end-port</b>] {<i>destination destination-host wildcard</i>   <b>any</b>   <b>host ip-address</b>} [{<b>eq</b>   <b>neq</b>   <b>gt</b>   <b>lt</b>} <i>dest-port</i>] [<b>range start-port end-port</b>] [<b>dscp code</b>] [<b>precedence value</b>] [<b>tos value</b>] [<b>log</b>   <b>log-verbose</b>]</p> <p><b>permit icmp</b> {<i>source source-wildcard</i>   <b>any</b>   <b>host ip-address</b>} {<i>destination destination-host wildcard</i>   <b>any</b>   <b>host ip-address</b>} [<b>msg icmp-msg</b>] [<b>dscp code</b>] [<b>precedence value</b>] [<b>tos value</b>] [<b>log</b>   <b>log-verbose</b>]</p>

**Procedure 38-4 Entering and Managing Extended IPv4 ACL Rules (continued)**

Step	Task	Command(s)
2.	In IPv4 ACL configuration command mode, optionally, create an extended IPv4 ACL deny rule entry.	<pre>deny {protocol-num   ip   ah   esp   gre} {source source-wildcard   any   host ip-address} {destination destination-host wildcard   any   host ip-address} [dscp code] [precedence value] [tos value] [log   log-verbose]  deny tcp {source source-wildcard   any   host ip-address} [{eq   neq   gt   lt} source-port] [range start-port end-port] {destination destination-host wildcard   any   host ip-address} [{eq   neq   gt   lt} dest-port] [range start-port end-port] [established] [dscp code] [precedence value] [tos value] [log   log-verbose]  deny udp {source source-wildcard   any   host ip-address} [{eq   neq   gt   lt} source-port] [range start-port end-port] {destination destination-host wildcard   any   host ip-address} [{eq   neq   gt   lt} dest-port] [range start-port end-port] [dscp code] [precedence value] [tos value] [log   log-verbose]  deny icmp {source source-wildcard   any   host ip-address} {destination destination-host wildcard   any   host ip-address} [msg icmp-msg] [dscp code] [precedence value] [tos value] [log   log-verbose]</pre>
3.	In IPv4 ACL configuration command mode, optionally, insert a new extended IPv4 ACL rule entry before the specified preexisting entry for this extended ACL. See the appropriate command syntax when entering a deny or permit rule to be inserted.	<pre>insert before entry {remark "text"   deny-syntax   permit-syntax}</pre>
4.	In IPv4 ACL configuration command mode, optionally, replace the specified extended IPv4 ACL entry with the specified new entry. See the appropriate command syntax when entering a deny or permit rule to be replaced.	<pre>replace entry {remark "text"   deny-syntax   permit-syntax}</pre>



Procedure 38-5 describes how to enter and manage extended IPv6 ACL rules.

### Procedure 38-5 Entering and Managing Extended IPv6 ACL Rules

Step	Task	Command(s)
1.	In IPv6 ACL configuration command mode, optionally, create an extended IPv6 ACL permit rule entry.	<pre> <b>permit</b> {<i>protocol-num</i>   <b>ipv6</b>   <b>ah</b>   <b>esp</b>   <b>gre</b>} {<i>source-address/length</i>   <b>any</b>   <b>host</b> <i>ip-address</i>} {<i>destination-address/length</i>   <b>any</b>   <b>host</b> <i>ip-address</i>} [<b>dscp</b> <i>code</i>] [<b>traffic-class</b> <i>value</i>] [<b>flow-label</b> <i>value</i>] [<b>log</b>   <b>log-verbose</b>] [<b>routing</b>] [<b>routing-type</b> <i>type</i>] [<b>mobility</b>] [<b>mobility-type</b> <i>type</i>]  <b>permit tcp</b> {<i>source-address/length</i>   <b>any</b>   <b>host</b> <i>ip-address</i>} [{<b>eq</b>   <b>neq</b>   <b>gt</b>   <b>lt</b>} <i>source-port</i>] [<b>range</b> <i>start-port end-port</i>] {<i>destination-address/length</i>   <b>any</b>   <b>host</b> <i>ip-address</i>} [{<b>eq</b>   <b>neq</b>   <b>gt</b>   <b>lt</b>} <i>dest-port</i>] [<b>range</b> <i>start-port end-port</i>] [<b>established</b>] [<b>dscp</b> <i>code</i>] [<b>traffic-class</b> <i>value</i>] [<b>flow-label</b> <i>value</i>] [<b>log</b>   <b>log-verbose</b>] [<b>routing</b>] [<b>routing-type</b> <i>type</i>] [<b>mobility</b>] [<b>mobility-type</b> <i>type</i>]  <b>permit udp</b> {<i>source-address/length</i>   <b>any</b>   <b>host</b> <i>ip-address</i>} [{<b>eq</b>   <b>neq</b>   <b>gt</b>   <b>lt</b>} <i>source-port</i>] [<b>range</b> <i>start-port end-port</i>] {<i>destination-address/length</i>   <b>any</b>   <b>host</b> <i>ip-address</i>} [{<b>eq</b>   <b>neq</b>   <b>gt</b>   <b>lt</b>} <i>dest-port</i>] [<b>range</b> <i>start-port end-port</i>] [<b>dscp</b> <i>code</i>] [<b>traffic-class</b> <i>value</i>] [<b>flow-label</b> <i>value</i>] [<b>log</b>   <b>log-verbose</b>] [<b>routing</b>] [<b>routing-type</b> <i>type</i>] [<b>mobility</b>] [<b>mobility-type</b> <i>type</i>]  <b>permit icmpv6</b> {<i>source-address/length</i>   <b>any</b>   <b>host</b> <i>ip-address</i>} {<i>destination-address/length</i>   <b>any</b>   <b>host</b> <i>ip-address</i>} [<i>icmpv6-type</i>   <i>icmpv6-code</i>]   <b>msg</b> <i>icmpv6-msg</i>] [<b>dscp</b> <i>code</i>] [<b>traffic-class</b> <i>value</i>] [<b>flow-label</b> <i>value</i>] [<b>log</b>   <b>log-verbose</b>] [<b>routing</b>] [<b>routing-type</b> <i>type</i>] [<b>mobility</b>] [<b>mobility-type</b> <i>type</i>] </pre>
2.		<pre> <b>deny</b> {<i>protocol-num</i>   <b>ipv6</b>   <b>ah</b>   <b>esp</b>   <b>gre</b>} {<i>source-address/length</i>   <b>any</b>   <b>host</b> <i>ip-address</i>} {<i>destination-address/length</i>   <b>any</b>   <b>host</b> <i>ip-address</i>} [<b>dscp</b> <i>code</i>] [<b>traffic-class</b> <i>value</i>] [<b>flow-label</b> <i>value</i>] [<b>log</b>   <b>log-verbose</b>] [<b>routing</b>] [<b>routing-type</b> <i>type</i>] [<b>mobility</b>] [<b>mobility-type</b> <i>type</i>]  <b>deny tcp</b> {<i>source-address/length</i>   <b>any</b>   <b>host</b> <i>ip-address</i>} [{<b>eq</b>   <b>neq</b>   <b>gt</b>   <b>lt</b>} <i>source-port</i>] [<b>range</b> <i>start-port end-port</i>] {<i>destination-address/length</i>   <b>any</b>   <b>host</b> <i>ip-address</i>} [{<b>eq</b>   <b>neq</b>   <b>gt</b>   <b>lt</b>} <i>dest-port</i>] [<b>range</b> <i>start-port end-port</i>] [<b>established</b>] [<b>dscp</b> <i>code</i>] [<b>traffic-class</b> <i>value</i>] [<b>flow-label</b> <i>value</i>] [<b>log</b>   <b>log-verbose</b>] [<b>routing</b>] [<b>routing-type</b> <i>type</i>] [<b>mobility</b>] [<b>mobility-type</b> <i>type</i>] </pre>

**Procedure 38-5 Entering and Managing Extended IPv6 ACL Rules (continued)**

Step	Task	Command(s)
3.	In IPv6 ACL configuration command mode, optionally, create an extended IPv6 ACL deny rule entry.	<pre>deny udp {source-address/length   any   host ip-address} [{eq   neq   gt   lt} source-port] [range start-port end-port] {destination-address/length   any   host ip-address} [{eq   neq   gt   lt} dest-port] [range start-port end-port] [dscp code] [traffic-class value] [flow-label value] [log   log-verbose] [routing] [routing-type type] [mobility] [mobility-type type]  deny icmpv6 {source-address/length   any   host ip-address} {destination-address/length   any   host ip-address} [icmpv6-type [icmpv6-code]   msg icmpv6-msg] [dscp code] [traffic-class value] [flow-label value] [log   log-verbose] [routing] [routing-type type] [mobility] [mobility-type type]</pre>
4.	In IPv6 ACL configuration command mode, optionally, insert a new extended IPv6 ACL rule entry before the specified preexisting entry for this extended ACL. See the appropriate command syntax when entering a deny or permit rule to be inserted.	<pre>insert before entry {remark "text"   deny-syntax   permit-syntax}</pre>
5.	In IPv6 ACL configuration command mode, optionally, replace the specified extended IPv6 ACL entry with the specified new entry. See the appropriate command syntax when entering a deny or permit rule to be replaced.	<pre>replace entry {remark "text"   deny-syntax   permit-syntax}</pre>

[Procedure 38-6](#) describes how to manage ACL rules.

**Procedure 38-6 Managing IPv4 and IPv6 ACL Rules**

Step	Task	Command(s)
1.	In IPv4 or IPv6 ACL configuration command mode, optionally, enable logging for the specified rule, the final implicit deny rule, or all rules.	<pre>log [entry] [implicit] [all]</pre>
2.	In IPv4 or IPv6 ACL configuration command mode, optionally, delete a preexisting ACL rule entry.	<pre>delete {entry   from entry to entry}</pre>
3.	In IPv4 or IPv6 ACL configuration command mode, optionally, move a preexisting ACL entry before the specified entry or range of entries.	<pre>move before entry from entry to entry</pre>
4.	In IPv4 or IPv6 ACL configuration command mode, optionally, enter a text comment as the next ACL entry.	<pre>remark "text"</pre>

[Procedure 38-7](#) describes how to apply and display ACLs.

### Procedure 38-7 Applying and Displaying ACLs

Step	Task	Command(s)
1.	In interface configuration command mode, apply an ACL to a routing interface specifying the whether the ACL applies to inbound or outbound frames.	<b>ipv4 access-group</b> { <i>access-list-number</i>   <i>name</i> } {in   out} <b>ipv6 access-group</b> <i>access-list-name</i> {in   out}
2.	In configuration command mode, apply an ACL to the host services for this device.	<b>ipv4 host-access</b> { <i>access-list-number</i>   <i>name</i> } <b>ipv6 host-access</b> <i>name</i>
3.	In any command mode, optionally, display ACL configuration.	<b>show access-lists</b> [ <i>access-list-number</i>   <i>name</i> ] [ <b>from</b> <i>start-range</i> <b>to</b> <i>end-range</i> ] [ <b>brief</b> ]
4.	In any command mode, optionally, display applied ACLs.	<b>show access-lists applied</b> [ <b>host</b>   <b>interfaces</b> [ <i>vlan</i>   <b>inbound</b>   <b>outbound</b>   <b>in-and-out</b> ]]
5.	In any command mode, optionally, clear ACL display counters.	<b>clear access-lists counters</b> [ { <i>access-list-number</i>   <i>name</i> }   <b>applied</b> [ <b>host</b>   <b>interfaces</b> [ <i>vlan</i> <i>vlan-id</i> ] [ <b>inbound</b>   <b>outbound</b>   <b>in-and-out</b> ] ]

## Terms and Definitions

[Table 38-1](#) lists terms and definitions used in this ACL configuration discussion.

**Table 38-1 ACL Configuration Terms and Definitions**

Term	Definition
Access Control List (ACL)	A container of permit, deny, and comment entries for the purpose of forwarding or not forwarding packets based upon one or more packet fields, such as source and destination IP address, and protocol.
entry	A member of an ACL that either permits or denies forwarding of the packet based upon one or more specified packet fields, or provides an ACL comment.
rule	An ACL entry that affects packet forwarding using a permit or deny entry.
standard ACL	An ACL for which forwarding decisions are made based only upon a source IP address.
extended ACL	An ACL for which forwarding decisions are made based upon the packet protocol, source and destination ip address, or host address, port matching in the case of the TCP or UDP protocols, as well as, optionally, a specified DSCP, ToS, or IP precedence value.

## Route-Map Manager Configuration

This document describes the route-map manager feature and its configuration on Extreme Networks 7100-Series devices.

For information about...	Refer to page...
<a href="#">Using Route-Map Manager in Your Network</a>	39-1
<a href="#">Implementing Route-Maps</a>	39-3
<a href="#">Route-Map Manager Overview</a>	39-4
<a href="#">Configuring Route-Map Manager</a>	39-8
<a href="#">Route-Map Manager Configuration Examples</a>	39-13
<a href="#">Terms and Definitions</a>	39-15

### Using Route-Map Manager in Your Network

The route-map manager supports four distinct types of route-maps:

- Redistribution route-maps provide for the filtering of routes redistributed from one routing domain to another via the OSPF protocol
- Policy based route-maps filter learned routes and support the calculation of the next hop forwarding decisions in a policy based routing context
- Filter route-maps provide for the denial of routes into the OSPF route table
- BGP route-maps provide for the permit and denial of BGP packets

A named route-map consists of a set of permit or deny entries. Entries are sequenced by unique sequence numbers per named route-map. A route-map can contain multiple route-map sequences. Route-map entries are not unlike the permit and deny statements in an ACL with one very important exception: unlike the ACL, all route-map entries must be successful for this route-map's action to occur.

Each route-map sequence may contain one or more match and set clauses. A match clause contains the criteria that determines whether the permit or deny action for this entry should be taken. All route-map entries for a given sequence must be successful for a route-map action to occur. If multiple sequences are configured, the first one that matches all entries will "pass" and return the set actions for that sequence. If a sequence does not pass, the next sequence is processed until a sequence in which all entries match is found. If no entries match for all sequences, then the route-map is not used.

A set clause defines the action for this route-map. Depending on the route-map type and permit/deny setting of the route-map sequence, zero or more set clauses are supported per route-map sequence.

## Policy Based Route-Maps

For policy based route-maps, if a match clause is configured, a match of the packet's source IP address against the contents of the specified ACL is required. A set entry specifies up to 5 next hop IP addresses for the forwarding of this packet.



**Note:** Policy Route-maps are limited to a single ACL with a single permit sequence (permit 10) in the match clause. The default, if none, is explicitly selected.

Policy based route-maps must be associated with an interface before route-mapping occurs. When assigning a route-map to an interface, the next hop load-policy behavior, which configures the algorithm used to select the next hop, and prioritization, which determines whether the priority based or routing table next hop is used, or whether the packet is dropped.

The route-map probe feature provides for the configuration of an ICMP probe to monitor next hops.

## Redistribution Route-Maps

For redistribution route-maps, if a match clause is configured, a match of the packet source IP address against either a specified VLAN or the contents of one or more specified ACLs is required. A configured set entry specifies a route tag, metric, metric increment or decrement, or metric type to be used for redistribution by the ACLs matched in this route-map.

Redistribution route-maps, with a set entry specifying a route tag, must be assigned to the **redistribute** command within the OSPF router configuration command mode, for redistribution based upon this route-map to occur.

## OSPF Filter Route-Maps

For OSPF filter route-maps, if a match clause is configured, a match on a deny route-map will deny the matched route from being installed into the OSPF route table based upon IP network address, next hop, source router-ID, outbound interface, OSPF tag, metric cost, or route-type.

OSPF filter route-maps must be assigned to the **distribution-list route-map in** command within OSPF configuration command mode for OSPF route table filtering to occur.

## BGP Route-Maps

For BGP permit route-maps, all match clauses within a sequence must match for set clauses to be performed. For BGP deny route-maps, all match clauses within a sequence must match for the packet to be dropped. There is an exception to the all match clauses rule: in the case of multiple match prefix entries, only a single prefix entry needs to match. BGP route-maps support match clauses for:

- Address Family Indicator (AFI) and Subsequent Address Family Indicator (SAFI) attributes
- AS-Path attribute
- Community name
- Extended-community name
- Prefix list
- Multi-Exit Discriminator (MED)
- Autonomous System (AS)

BGP route-maps support set clauses for:

- Autonomous System (AS)
- Maximum length of the AS path attribute

- Community name
- Extended community attributes:
  - IP route target
  - AS and 4-octet AS route target
  - IP site of origin
  - AS and 4-octet AS site of origin
  - OSPF domain and router ID
  - OSPF route type
- Local preference
- Multi-Exit Discriminator (MED)
- IP next hop
- Origin
- Local Outbound Rate Filtering (ORF) association
- Weight
- Flap table

## Implementing Route-Maps

### Implementing a Policy Based Route-Map

To implement a policy based route-map:

- Create a policy based route-map and one or more entries for this route map
- For each sequence in the route-map, optionally configure match clauses to filter the packet based upon the specification of an ACL per match clause
- For each sequence in the route-map, optionally configure a set clause specifying up to 5 next hops per command line (system maximum of 128)
- Optionally configure the route-map probe feature to monitor each specified next hop in the route-map
- Assign the configured route-map to the interface for which policy-based routing is to be performed (a route-map can be assigned to multiple interfaces)

### Implementing a Redistribution Route-Map

To implement a redistribution route-map:

- Create a redistribution route-map and one or more entries for this route-map
- For each sequence in the route-map, optionally configure match clauses to filter the packet source IP address based upon the specification of up to five ACLs per match clause
- For each sequence in the route-map, optionally configure match clauses to filter the packet source IP address based upon a specified interface
- For each sequence in the route-map, optionally configure match clauses to filter the packet based upon route cost or a route cost range.

- For each sequence in the route-map, optionally configure a set clause containing an OSPF route tag or range of route tags for this route-map
- In router configuration command mode, assign the route-map to the redistribute feature

## Implementing an OSPF Filter Route-Map

To implement an OSPF filter route-map:

- Create a filter route-map and one or more entries for this route-map
- For each sequence in the route-map, optionally configure match clauses to filter routes for this OSPF route table
- In OSPF router configuration command mode, apply the route map filter using the `distribute-list route-map` in command

## Implementing a BGP Route-Map

To implement a BGP route-map:

- Create a BGP route-map and one or more entries for this route-map.
- For each sequence in the route-map, configure match clauses. See “[BGP Route-Maps](#)” on page 39-2 for a listing of supported match clauses.
- For each sequence in the route-map, optionally configure set clauses. See “[BGP Route-Maps](#)” on page 39-2 for a listing of supported set clauses.

## Route-Map Manager Overview

This section provides an overview of route-map manager configuration.

### Creating a Route-Map

When creating a route-map, specify:

- Whether it is a policy based, redistribution, or filter route-map
- The name of the route-map using up to 32 alpha-numeric characters
- Whether this sequence is a permit or deny (defaults to permit)
- A sequence number for this entry (defaults to 10)

Currently, up to 100 each of filter, redistribution, and policy route-maps are permitted.

Multiple sequences can be input for a single named route-map. Configuring a route-map sequence places you in route-map configuration command mode for the configuration of route-map match and set clauses. The system-wide maximum number of both match and set route-map clauses is 1000.

Policy route-maps must have at least one IP address match clause and at least one next hop clause. An ACL that has not yet been created can be specified in an IP address match clause. If a route-map is applied to an interface, any ACLs that have not been created will be ignored. Policy based route-maps must be assigned to an interface using the **ip policy route-map** command in interface configuration mode.

Redistribution route-maps must be associated with the redistribution of OSPF routes within the OSPF routing protocol using the **redistribute** command in OSPF router configuration command mode.

Filter route-maps must be associated with the filtering of OSPF routes from the OSPF route table (FIB) using the **distribute-list route-map in** command.

Use the **route-map policy** command in configuration command mode to create a policy based route-map.

Use the **route-map redistribution** command in configuration command mode to create a redistribution route-map.

Use the **route-map filter** command in configuration command mode to create an OSPF filter route-map.

## Configuring Match and Set Clauses

Upon entering a route-map sequence, you are placed in route-map configuration command mode. Match and set clauses are configured in route-map configuration command mode.

A route-map sequence's match clause specifies the criterion that determines whether the action for this route-map will occur. The following types of match clauses are supported:

### Redistribution Match Clauses

- A match clause that matches packets egressing on this interface with the statements in up to five specified ACLs. Multiple clauses may be used. At least one of the ACLs in each clause must match the packet in order for the route-map to redirect the packet. The only limit on the number of ACLs supported is the system limit of 1000 route-map clauses. Use the **match ip address** command in redistribution route-map configuration command mode to specify up to five ACLs for this match clause.
- An interface match clause that matches the source IP address of a packet egressing on this interface against a specified VLAN interface. Use the **match interface** command in redistribution route-map configuration command mode to specify a VLAN interface for this match clause.
- A metric match clause that matches the specified or a range of cost against the route cost specified in the packet. Use the **match metric** command in redistribution route-map configuration command mode to specify the metric cost for this match clause.
- An OSPF tag match clause that matches the specified OSPF tag or range of tags against the OSPF tag ID specified in the packet. Use the **match tag** command in redistribution route-map configuration command mode to specify the OSPF tag ID for this match clause.

### Policy Match Clauses

An IP address match clause that matches the source IP address of a packet egressing on this interface with the statements in a single ACL. The IP address match clause can be entered for both a policy based route-map and a redistribution route-map.

Use the **match ip address** command in policy-based route-map configuration command mode to specify a single ACL to be associated with this match clause. The ACL must match the packet in order for the route-map to redirect the packet.

### OSPF Filter Match Clauses

- An IP match clause that matches a route network address, next hop or source router ID against the route to be entered into the OSPF routing table. Use the **match ip** command in filter-based route-map configuration command mode to specify a single ACL to be associated with this match clause. The ACL must match the packet in order for the route-map to redirect the packet.



- An interface match clause that matches the outgoing interface of the route to be installed in the OSPF routing table. Use the **match interface** command in filter-based route-map configuration command mode to specify an outgoing interface for this match clause.
- A OSPF tag match clause that matches the OSPF tag for this route. Use the **match tag** command in filter-based route-map configuration command mode to specify an OSPF tag or range of tags for this match clause.
- A metric match clause that matches the OSPF cost metric for this route. Use the **match metric** command in filter-based route-map configuration command mode to specify an OSPF route cost metric or range of cost metrics for this match clause.
- A route-type match clause that matches the internal or external route type for this route. Use the **match route-type** command in filter-based route-map configuration command mode to specify an OSPF route-type for this match clause.

There can be multiple match clauses associated with a single route-map sequence.

A route-map sequence's set clause determines the action the route-map will take when a successful match for this sequence occurs. The action configurable for a set clause depends upon the route-map type. For a policy based route-map, the set clause specifies one or more next hops for this route. For the redistribution route-map, the set clause specifies an OSPF route tag for this route.

## Policy Based Set Clauses

Policy based set clauses determine the next hop for this route if a match clause for this route-map sequence is successful.

Use the **set next-hop** command in policy based route-map configuration mode to specify the next hop(s) available for this route-map's action.

## Route-Map Probe

The route-map manager supports the assigning of an ICMP probe to monitor a next hop IP address. Tracked object manager uses the route-map facility to monitor the IP address, but you do not assign the ICMP probe to a specific route-map. If a next hop IP address is declared down, it is removed from the next hop selection process for all route-maps specifying this address as a next hop, until it is declared up again. The assigned ICMP probe will ping port 0 of the specified IPv4 or IPv6 address.

A route-map probe entry is configurable for each configured next hop address. Currently a combination of up to 128 standard next hop addresses are configurable on a system. If the same next hop is referenced in multiple route-maps, only a single route-map probe instance is created.

See [Chapter 12, Tracked Object Manager Configuration](#) for tracked object manager details.

Use the **route-map probe** command in router configuration mode to assign an ICMP probe to monitor the specified next hop IP address. A predefined policy based routing ICMP probe named **\$pbr\_default** can be used, or you can create a probe, using the **probe** command. Predefined ICMP probes can not be specified by name. Use the **default** keyword when configuring the default route-map probe.

This example shows how to create the ICMP probe **ICMP-PBR** and assign it to a route-map probe to monitor next hop IP address **101.10.1.252**. The fail detection count is set to **5** attempts, and the fail detection interval is set to **5** seconds. The assigned session is displayed:

```
7100 Chassis(su-config)->probe ICMP-PBR icmp
7100 Chassis(su-config-probe)->faildetect count 5 interval 5
7100 Chassis(su-config-probe)->inervice
7100 Chassis(su-config-probe)->exit
```

```

7100 Chassis(su-config)->route-map probe 101.10.1.252 probe-name ICMP-PBR
7100 Chassis(su-config)->show probe sessions

Client Codes: P-policy based routing, V-VRRP
               T-tracked object probe
...
Probe: ICMP-PBR, icmp
IP Address                               Port  Status   StChngs Last Change  Clients
-----
101.10.1.252                             0     Up       1         0h0m30s    P
Displayed 1 sessions
...
7100 Chassis(su-config)->

```

## The Redistribution Match Clauses

The redistribution route-map entry allows the specifying of both IP address and interface match clauses. A single ACL can be configured in an IP address match clause. A single interface can be configured for an interface match clause.

## The Redistribution Set Clause

The redistribution set clause determines the OSPF route tag, metric cost, along with the ability to increment or decrement the current metric cost, and metric type for this route if a match clause for this route-map entry is successful.

Use the **set tag** command in redistribution route-map configuration command mode to specify the OSPF route to be used for redistributing non-OSPF routes that match for this route-map.

Use the **set metric** command in redistribution route-map configuration command mode to specify the metric cost of routes that match for this route-map. Use the **set metric increment** command to increase the current metric cost or **set metric decrement** command to decrement the current metric cost of routes that match for this route-map.

OSPF route tag is a 32-bit numeric value that is attached to redistributed routes into OSPF. The route tag is not used by OSPF, but can be used by other routers for making policy decisions. OSPF route tags are displayed in the **show ip ospf database external** command. See the *Extreme Networks 7100-Series CLI Reference* for command details.

## Assigning a Policy Route-Map to an Interface

Route-map filtering does not occur until the configured route-map is assigned to an interface. Once assigned to an interface the route-map is operational.

Next hop load-policy is hard coded to first-available next hop from the list of next hops.

Priority is hard coded to prioritize use of the priority-based routing by using the priority-based routing next hop, but if it is unavailable, the route table is used.

Use the **ip policy route-map** command in interface configuration command mode to assign a route-map to an interface.

## Configuring Route-Map Manager

This section provides details for the configuration of route-map manager on the 7100-Series products.

[Table 39-1](#) lists route-map manager parameters and their default values.

**Table 39-1 Default Route-Map Manager Parameters**

Parameter	Description	Default Value
entry	A route-map's sequenced container for match and set clauses specifying a permit or deny behavior.	permit
sequence number	A numeric value specifying the ordering of route-map entries.	10
next hop priority	Specifies whether the priority based lookup or the routing table lookup will be used to select the next hop. Non-Configurable.	priority based lookup, then route table lookup
next hop load-policy	Specifies the algorithm that will be used to select the next hop. Non-Configurable.	first-available

[Procedure 39-1](#) describes how to configure a policy based route-map.

**Procedure 39-1 Configuring a Policy Based Route-Map**

Step	Task	Command(s)
1.	In configuration command mode, create a policy based route map, optionally specifying whether this entry is a permit or deny, and the sequence number for this entry.  This command provides access to policy based route-map configuration command mode. Use this command to create multiple entries if required.	<b>route-map policy</b> <i>name</i> [ <b>permit</b>   <b>deny</b> ] [ <i>sequence-number</i> ]
2.	In policy based route-map configuration command mode, specify one or more match clauses for this route-map, specifying up to five ACLs that will be matched against the packet source IP address. Though not necessary, it is recommended that all ACLs be configured before assigning them to an IP address match clause.	<b>match ip address</b> <i>access-list</i>
3.	In policy based route-map configuration command mode, specify a set clause containing up to five next hop IP addresses for this route-map. One or more of these commands can be specified.	<b>set next-hop</b> { <i>next-hop1</i> } [ <i>next-hop2</i> .... <i>next-hop5</i> ]
4.	Optionally, in configuration command mode, configure the route-map probe feature to monitor the configured next hops.	<b>route-map probe</b> <i>ip-address</i> <b>probe-name</b> { <i>name</i>   <b>default</b> }
5.	In interface configuration command mode, assign the configured route-map to the interface.	<b>ip policy route-map</b> <i>name</i>

[Procedure 39-2](#) describes how to configure a redistribution route-map.

### Procedure 39-2 Configuring a Redistribution Route-Map

Step	Task	Command(s)
1.	In configuration command mode, create a redistribution route map, optionally specifying whether this entry is a permit or deny, and the sequence number for this entry.  This command provides access to redistribution route-map configuration command mode. Use this command to create multiple entries if required.	<b>route-map redistribution</b> <i>name</i> [ <b>permit</b>   <b>deny</b> ] [ <i>sequence-number</i> ]
2.	In redistribution route-map configuration command mode, specify one or more match clauses for this route-map, specifying up to five ACLs that will be matched against the packet source IP address.	<b>match ip address</b> <i>access-list</i>
3.	In redistribution route-map configuration command mode, specify a VLAN interface to match a packet source IP address against.	<b>match interface</b> { <i>vlan vlan</i>   <i>string</i> }
4.	In redistribution route-map configuration command mode, specify one or a range of metric costs that will be matched against the packet metric cost.	<b>match metric</b> { <i>cost</i>   <b>range</b> <i>min-cost</i> <i>max-cost</i> }
5.	In redistribution route-map configuration command mode, specify an OSPF tag ID or range of IDs that will be matched against the packet OSPF tag ID.	<b>match tag</b> { <i>tag-id</i>   <b>range</b> <i>min-tag-id</i> <i>max-tag-id</i> }
6.	In redistribution route-map configuration command mode, specify a set clause containing an OSPF route tag for this route-map.	<b>set tag</b> <i>tag</i>
7.	In redistribution route-map configuration command mode, specify a set clause containing a metric cost for this route-map. A single metric cost can be configured per sequence.	<b>set metric</b> <i>cost</i>
8.	In redistribution route-map configuration command mode, specify a set clause containing the amount to decrement the current metric cost for this route-map. A single metric decrement can be configured per sequence.	<b>set metric decrement</b> <i>cost</i>
9.	In redistribution route-map configuration command mode, specify a set clause containing the amount to increment the current metric cost for this route-map. A single metric increment can be configured per sequence.	<b>set metric increment</b> <i>cost</i>
10.	In redistribution route-map configuration command mode, specify a set clause containing the OSPF metric type to be used when redistributing a source packet matched by this route-map. A single metric type can be configured per sequence.	<b>set metric-type</b> { <b>type-1</b>   <b>type-2</b> }

**Procedure 39-2 Configuring a Redistribution Route-Map (continued)**

Step	Task	Command(s)
11.	In OSPF router configuration mode, assign this route-map to the redistribute command.	<b>redistribute</b> {rip   static   connected} [route-map name] [metric metric value] [metric-type type-value] [tag tag]

[Procedure 39-3](#) describes how to configure an OSPF filter route-map.

**Procedure 39-3 Configuring a Filter Route-Map**

Step	Task	Command(s)
1.	In configuration command mode, create an OSPF filter route map, optionally specifying whether this entry is a permit or deny, and the sequence number for this entry.  This command provides access to filter route-map configuration command mode. Use this command to create multiple entries if required.	<b>route-map filter</b> name [permit   deny] [sequence-number]
2.	In filter route-map configuration command mode, specify one or more IP network address, next hop, or source router ID match clauses for this route-map, specifying up to five ACLs that will be matched against specified IP type. <ul style="list-style-type: none"> <li>• <b>address</b> - network address</li> <li>• <b>next-hop</b> - next hop</li> <li>• <b>route-source</b> - source router ID</li> </ul>	<b>match ip</b> {address   next-hop   route-source} access-list
3.	In filter route-map configuration command mode, specify one or more outbound interface match clauses that will be matched against the route outbound interface.	<b>match interface</b> {interface-name   alias}
4.	In filter route-map configuration command mode, specify one or more OSPF tag match clauses that will be matched against the route OSPF tag or a range of OSPF tags.	<b>match tag</b> {tag   range min-tag max-tag}
5.	In filter route-map configuration command mode, specify one or more OSPF cost metric match clauses that will be matched against the route metric cost or a range of metric cost values.	<b>match metric</b> {cost   range min-cost max-cost}
6.	In filter route-map configuration command mode, specify one or more OSPF route type match clauses that will be matched against the route's route type. <ul style="list-style-type: none"> <li>• <b>internal</b> - Internal route type</li> <li>• <b>external-t1</b> - External route type 1</li> <li>• <b>external-t2</b> - External route type 2</li> <li>• <b>nssa-external</b> - External NSSA route type</li> </ul>	<b>match route-type</b> {internal   external-t1   external-t2   nssa-external}
7.	In OSPF router configuration command mode, apply the filter route-map to the OSPF distribution-list.	<b>distribute-list route-map</b> name in

[Procedure 39-4](#) describes how to configure a BGP route-map.

#### Procedure 39-4 Configuring a BGP Route-Map

Step	Task	Command(s)
1.	In configuration command mode, create a BGP route map, optionally specifying whether this entry is a permit or deny, and the sequence number for this entry.  This command provides access to BGP route-map configuration command mode. Use this command to create multiple entries if required.	<b>route-map</b> <i>bgp name</i> [ <b>permit</b>   <b>deny</b> ] [ <i>sequence-number</i> ]
2.	In BGP route-map configuration command mode, configure a match clause to match a packet against its IPv4 Address Family Indicator (AFI) attribute.	<b>match</b> <i>afi ipv4</i>
3.	In BGP route-map configuration command mode, configure a match clause to match a packet against its Subsequent Address Family Indicator (SAFI) attribute, specifying whether the attribute is unicast or multicast.	<b>match</b> <i>safi {unicast   multicast}</i>
4.	In BGP route-map configuration command mode, configure a match clause to match a packet against its AS path attribute.	<b>match</b> <i>as-path as-path-string</i>
5.	In BGP route-map configuration command mode, configure a match clause to match a packet against the specified community name.	<b>match</b> <i>community name</i>
6.	In BGP route-map configuration command mode, configure a match clause to match a packet against the specified extended community.	<b>match</b> <i>extended-community name</i>
7.	In BGP route-map configuration command mode, configure a match clause to match a packet against the specified prefix list. Multiple prefix-list match entries are allowed.	<b>match</b> <i>prefix-list prefix-list</i>
8.	In BGP route-map configuration command mode, configure a match clause to match a packet against the specified MED value.	<b>match</b> <i>med value</i>
9.	In BGP route-map configuration command mode, configure a match clause to specify the number of times to prepend the AS number of this router to the AS path for this route map context.	<b>set</b> <i>as num</i>
10.	In BGP route-map configuration command mode, configure a match clause to set a maximum length of the AS path attribute allowed when all match clauses match for this route map.	<b>set</b> <i>as-path-limit limit</i>
11.	In BGP route-map configuration command mode, configure a match clause to set the community when all match clauses match for this route map.	<b>set</b> <i>community {as:community   defined-community}</i>

**Procedure 39-4 Configuring a BGP Route-Map (continued)**

Step	Task	Command(s)
12.	In BGP route-map configuration command mode, configure a match clause to specify an action for an extended community IP route target when all match clauses match for this route map.	<b>set extended-community ip-route-target</b> <i>set-value</i> {remove-all   remove-specific   set-specific   remove-all-and-set}
13.	In BGP route-map configuration command mode, configure a match clause to specify an action for an extended community AS route target when all match clauses match for this route map.	<b>set extended-community as-route-target</b> <i>set-value</i> {remove-all   remove-specific   set-specific   remove-all-and-set}
14.	In BGP route-map configuration command mode, configure a match clause to specify an action for an extended community IP site of origin when all match clauses match for this route map.	<b>set extended-community ip-site-of-origin</b> <i>set-value</i> {remove-all   remove-specific   set-specific   remove-all-and-set}
15.	In BGP route-map configuration command mode, configure a match clause to specify an action for an extended community AS site of origin when all match clauses match for this route map.	<b>set extended-community as-site-of-origin</b> <i>set-value</i> {remove-all   remove-specific   set-specific   remove-all-and-set}
16.	In BGP route-map configuration command mode, configure a match clause to specify an action for an extended community AS4 route target when all match clauses match for this route map.	<b>set extended-community as4-route-target</b> <i>set-value</i> {remove-all   remove-specific   set-specific   remove-all-and-set}
17.	In BGP route-map configuration command mode, configure a match clause to specify an action for an extended community AS4 site of origin when all match clauses match for this route map.	<b>set extended-community as4-site-of-origin</b> <i>set-value</i> {remove-all   remove-specific   set-specific   remove-all-and-set}
18.	In BGP route-map configuration command mode, configure a match clause to specify an action for an extended community OSPF domain ID when all match clauses match for this route map.	<b>set extended-community ospf-domain-id</b> <i>set-value</i> {remove-all   remove-specific   set-specific   remove-all-and-set}
19.	In BGP route-map configuration command mode, configure a match clause to specify an action for an extended community OSPF domain ID when all match clauses match for this route map.	<b>set extended-community ospf-router-id</b> <i>set-value</i> {remove-all   remove-specific   set-specific   remove-all-and-set}
20.	In BGP route-map configuration command mode, configure a match clause to specify an action for an extended community OSPF route type when all match clauses match for this route map.	<b>set extended-community ospf-route-type</b> <i>area route-type type</i> [type2-metric] {remove-all   remove-specific   set-specific   remove-all-and-set}
21.	In BGP route-map configuration command mode, configure a match clause to specify the local preference to be set when all match clauses in the route map match.	<b>set local-preference</b> <i>value</i>

**Procedure 39-4 Configuring a BGP Route-Map (continued)**

Step	Task	Command(s)
22.	In BGP route-map configuration command mode, configure a match clause to specify the MED to be set when all match clauses in the route map match.	<b>set med <i>value</i></b>
23.	In BGP route-map configuration command mode, configure a match clause to specify the next hop IP address to be set when all match clauses in the route map match.	<b>set ip next-hop <i>ip-address</i></b>
24.	In BGP route-map configuration command mode, configure a match clause to specify the origin code to be set when all match clauses in the route map match.	<b>set origin <i>code</i></b>
25.	In BGP route-map configuration command mode, configure a match clause to set local ORF association when all match clauses in the route map match.	<b>set orf-association local</b>
26.	In BGP route-map configuration command mode, configure a match clause to specify the weight to be set when all match clauses in the route map match.	<b>set weight <i>value</i></b>
27.	In BGP route-map configuration command mode, configure a match clause to specify the flap table to be set when all match clauses in the inbound route map match.	<b>set flap-table <i>name</i></b>

[Table 39-2](#) describes how to display route-map manager information. Display commands can be entered in any command mode.

**Table 39-2 Displaying Route-Map Manager Information and Statistics**

Task	Command
To display configured route-maps:	<b>show route-map [<i>name</i>] [<i>brief</i>] [<i>probe</i>]</b>
To display the policy applied to a routing interface:	<b>show ip policy</b>

## Route-Map Manager Configuration Examples

This section presents a route-map manager configuration examples for a policy based and a redistribution route-map.

### Policy Based Route-Map Example

The following example:

- Creates a policy based route-map name **rmP1** that filters IP packets with source addresses on the **60.10.0.0** subnet destined for hosts **50.10.0.1-2**.
- Packets that pass this filter will be routed using one of three next hops: **30.10.0.10**, **30.10.0.20**, or **30.10.0.30**.
- The route-map probe feature is configured to monitor these three next hops for availability using the default policy based routing probe **\$pbr\_default**.



- The route-map is assigned to VLAN 110.

```
7100 Chassis(rw)->configure
7100 Chassis(rw-config)->ip access-list extended 101
7100 Chassis(rw-cfg-ext-acl)->permit ip 60.10.0.0 0.0.255.255 host 50.10.0.1
7100 Chassis(rw-cfg-ext-acl)->permit ip 60.10.0.0 0.0.255.255 host 50.10.0.2
7100 Chassis(rw-cfg-ext-acl)->deny ip any any
7100 Chassis(rw-cfg-ext-acl)->show access-lists 101
Extended IP access list 101 (4 entries)
  1 permit ip 60.10.0.0 0.0.255.255 host 50.10.0.1
  2 permit ip 60.10.0.0 0.0.255.255 host 50.10.0.2
  3 deny ip any any
  -- implicit deny all --
7100 Chassis(rw-cfg-ext-acl)->exit
7100 Chassis(rw-config)->route-map policy rmP1 permit 10
7100 Chassis(rw-config-route-map-pbr)->match ip address 101
7100 Chassis(rw-config-route-map-pbr)->set next-hop 30.10.0.10 30.10.0.20
30.10.0.30
7100 Chassis(rw-config-route-map-pbr)->exit
7100 Chassis(rw-config)->show route-map rmP1
route-map policy rmP1 permit 10
  match ip address 101
  set next-hop 30.10.0.10 30.10.0.20 30.10.0.30
Policy matches: 0 packets
7100 Chassis(rw-config)->route-map probe 30.10.0.10 default
7100 Chassis(rw-config)->route-map probe 30.10.0.20 default
7100 Chassis(rw-config)->route-map probe 30.10.0.30 default
7100 Chassis(rw-config)->interface vlan 110
7100 Chassis(rw-config-intf-vlan.0.110)->ip policy route-map rmP1
```

## Redistribution Route-Map Example

The following example:

- Creates a redistribution route-map named **rmR1** for the redistribution of RIP routes with a permit entry, sequence 10 that filters IP packets with source addresses on the 40.0.0.0 and 40.0.10.0 subnets
- Packets that pass the filter have the OSPF route tag set to 65432
- Redistribute in OSPF router 1 is assigned the rmR1 route-map

```
7100 Chassis(rw)->configure
7100 Chassis(rw-config)->ip access-list standard OSPF
7100 Chassis(rw-cfg-std-acl)->permit 40.0.0.0 0.0.0.255
7100 Chassis(rw-cfg-std-acl)->permit 40.0.10.0 0.0.0.255
7100 Chassis(rw-cfg-std-acl)->show access-lists OSPF
Standard IP access list OSPF (3 entries)
  1 permit 40.0.0.0 0.0.0.255
  2 permit 40.0.10.0 0.0.0.255
```

```

-- implicit deny all --
7100 Chassis(rw-cfg-std-acl)->exit
7100 Chassis(rw-config)->route-map redistribution rmR1 permit 10
7100 Chassis(rw-config-route-map)->match ip address OSPF
7100 Chassis(rw-config-route-map)->set tag 65432
7100 Chassis(rw-config-route-map)->exit
7100 Chassis(rw-config)->show route-map rmR1
    route-map redistribution rmR1 permit 10
    match ip address OSPF
    set tag 65432
7100 Chassis(rw-config)->router ospf 1
7100 Chassis(rw-config-ospf-1)->redistribute rip route-map rmR1
7100 Chassis(rw-config-ospf-1)->exit
7100 Chassis(rw-config)->

```

## BGP Route-Map Example

The following BGP route-map example:

- Creates a BGP route-map named **bgprm1** as a permit entry
- Specifies that the packet prefix should match a prefix listed in prefix list **pxlist1**
- Specifies that the packet AS path string should match **^20313.\*13\$**
- Specifies the setting of the IP next hop to **152.50.25.10** as the action if all match clauses match

```

7100 Chassis(su)->configure
7100 Chassis(su-config)->route-map bgp bgprm1 permit
7100 Chassis(su-config-route-map-bgp)->match prefix-list permit100
7100 Chassis(su-config-route-map-bgp)->match prefix-list pxlist1
7100 Chassis(su-config-route-map-bgp)->match as-path ^20313.*13$
7100 Chassis(su-config-route-map-bgp)->set ip next-hop 152.50.25.10
7100 Chassis(su-config-route-map-bgp)->

```

## Terms and Definitions

[Table 39-3](#) lists terms and definitions used in this route-map manager configuration discussion.

**Table 39-3 Route-Map Manager Terms and Definitions**

Term	Definition
entry	A logical container within the named route-map that specifies a permit or deny behavior for the configured match and set clauses it contains.
filter route-map	A route filtering container that provides for the denial of routes into the OSPF route table.
load-policy	The ability to configure the algorithm that will be used for the next hop selection for this route-map.
match clause	A clause that specifies the criteria for filtering routes for a route-map.

**Table 39-3 Route-Map Manager Terms and Definitions (continued)**

<b>Term</b>	<b>Definition</b>
route-map probe	A tracked object manager object of protocol type ICMP that tracks the availability of a next hop IP address, by actively pinging the address.
policy route-map	A route filtering container that permits or denies routes based upon an ACL entry match, optionally allowing for the specification of up to five next hops for routes that pass the filter.
priority	The ability to configure whether the priority route lookup or the route table lookup will determine the next hop for this route.
redistribution route-map	A route filtering container that permits or denies routes based upon an ACL entry match for purposes of redistribution over the OSPF protocol
set clause	A clause that specifies the action that will occur for routes matched by the route-map match clause.

## Quality of Service (QoS) Configuration

This chapter describes the QoS feature as it is implemented on the Extreme Networks 7100-Series devices.

For information about...	Refer to page...
<a href="#">Using Quality of Service in Your Network</a>	40-1
<a href="#">Implementing Quality of Service</a>	40-2
<a href="#">Quality of Service Overview</a>	40-2
<a href="#">Understanding QoS Configuration on the 7100-Series</a>	40-8
<a href="#">The QoS CLI Command Flow</a>	40-17
<a href="#">Terms and Definitions</a>	40-18

### Using Quality of Service in Your Network

Quality of Service (QoS) is:

- A mechanism for the management of bandwidth
- The ability to give preferential treatment to some packets over others
- Based upon packet classification and forwarding treatment

You configure packet preference and forwarding treatment based upon a flow's sensitivity to delay, delay variation (jitter), bandwidth, availability and packet drop. QoS uses packet priority, in conjunction with queue treatment configuration, to determine the interface's inbound and forwarding behavior for this packet. Packet preference and forwarding treatment for a given flow can be applied to roles configured in Extreme Networks policy.

Without QoS, all packets are treated as though the delivery requirements and characteristics of any given packet are equal to any other packet. In other words, non-QoS packet delivery is not able to take into account application sensitivity to packet delay, jitter, amount of bandwidth required, packet loss, or availability requirements of the flow. QoS provides management mechanisms for these flow characteristics.

QoS achieves its bandwidth management capabilities by:

- Setting priorities that define traffic handling
- Dedicating bandwidth and prioritizing queuing for specific applications, and reducing packet transmission delay and jitter
- Managing congestion by shifting packet loss to applications that can tolerate it

## Implementing Quality of Service

QoS determines how a flow will be treated as it transits the link. To determine how a flow should be treated, you must first understand the characteristics of the flows on your network, and secondly, you must identify these flows in a way that QoS can recognize. In this sense, QoS is the third step in a three step process. The three steps Extreme Networks recommends for configuring QoS are:

- Understand your network flows.
- Associate the flows on your network with a well defined role using Extreme Networks policy. See [Chapter 22, Policy Configuration](#) for policy configuration details.
- Configure the appropriate link behavior for that role by associating the role with a QoS configuration.

## Quality of Service Overview

QoS is all about managing the bandwidth in a manner that aligns the delivery characteristics of a given flow with the available port resources. In a QoS context, a flow is a stream of IP packets that are classified with the same class of service as it transits the interface. QoS manages bandwidth for each flow by taking advantage of its ability to:

- Assign different priority levels to different packet flows
- Mark or re-mark the packet priority at port ingress with a Type of Service
- Sort flows by transit queue such that a higher priority queue gets preferential access to bandwidth during packet forwarding
- Limit the amount of bandwidth available to a given flow by either dropping (rate limiting) or buffering (rate shaping) packets in excess of configured limits

These QoS abilities collectively make up a Class of Service (CoS). The remainder of this section will briefly describe CoS and its components.

## Class of Service (CoS)

You implement QoS features in a Class of Service (CoS). How the firmware treats a packet as it transits the link depends upon the priority and forwarding treatments configured in the CoS. Up to 256 unique CoS entries can be configured. CoS entries 0–7 are configured by default with an 802.1p priority assigned and default forwarding treatment. For purposes of backward compatibility, CoS entries 0–7 cannot be removed. CoS entries 8–255 can be configured for the following services:

- 802.1p priority
- IP ToS rewrite value
- Priority Transmit Queue (TxQ) with configurable forwarding behavior
- In-bound (IRL) rate limiter per transmit queue
- Outbound rate shaper per transmit queue

The CoS configuration for each service can be easily viewed using the CoS setting tables. Ports are bundled into port groups with the group assigned to a CoS, significantly cutting down on operational overhead and complexity.

## CoS Priority and ToS Rewrite

The two parameters configurable for CoS priority are 802.1p and Type of Service (ToS). Each CoS can be mapped to an 802.1p priority and a ToS rewrite value. The 802.1p parameter is:

- A subset of ToS with values 0–7 (upper 3 bits of the 8 bit ToS field)
- Supported in both layer 2 and layer 3

The ToS parameter is:

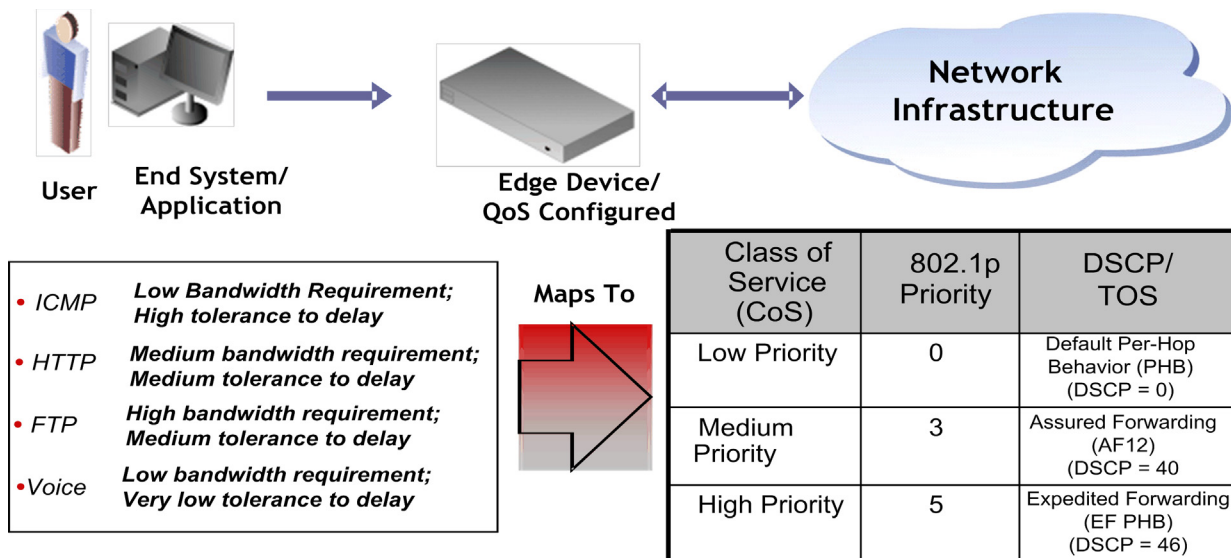
- An 8-bit field with values 0–255
- Supported in layer 3 only
- Also referred to as the Differentiated Services Code Point (DSCP) when limited to the lower 5 bits of the field

Figure 40-1 displays the relationship between your application, priority level, 802.1p, and ToS assignments (shown here using DSCP terminology).

QoS priority/ToS configuration:

- Derives its characteristic requirements from the end-system application
- Is configured on the edge device the application is connected to
- Is propagated through the network in the protocol packet header

**Figure 40-1 Assigning and Marking Traffic with a Priority**



The ICMP protocol, used for error messaging, has a low bandwidth requirement, with a high tolerance for delay and jitter, and is appropriate for a low priority setting. HTTP and FTP protocols, used respectively for browser-generated and file transfer traffic, have a medium to high bandwidth requirement, with a medium to high tolerance for delay and jitter, and are appropriate for a medium priority level. Voice (VoIP), used for voice calls, has a low bandwidth requirement, but is very sensitive to delay and jitter and is appropriate for a high priority level.

See RFC 1349 for further details on ToS. See RFCs 2474 and 2475 for further details on DSCP.

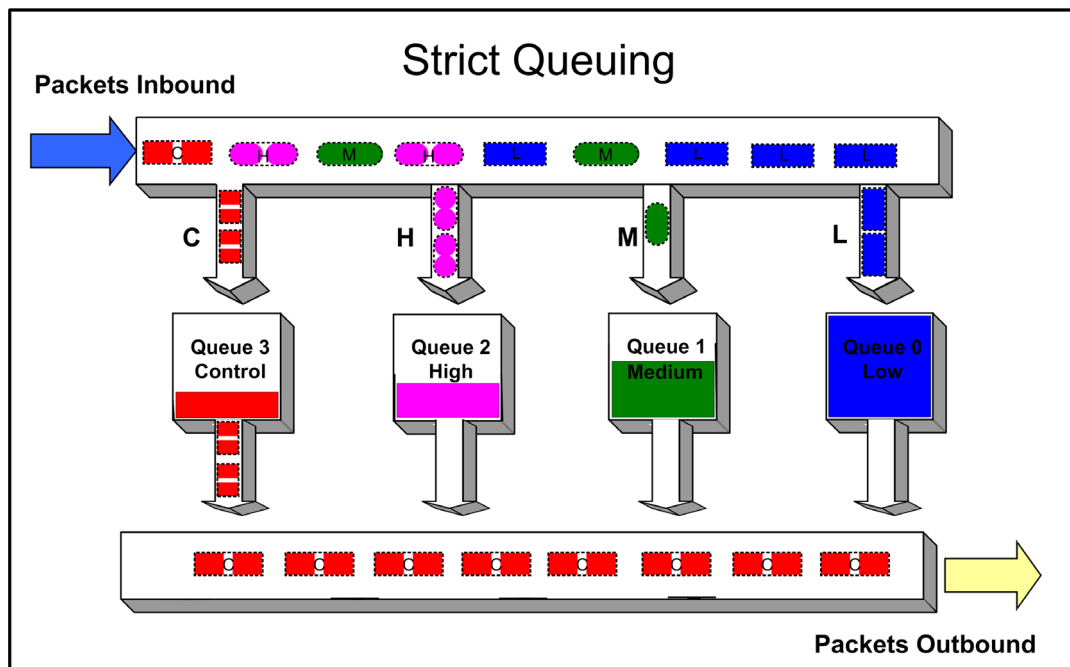
## Preferential Queue Treatment for Packet Forwarding

There are three types of preferential queue treatments for packet forwarding: strict priority, weighted fair, and hybrid.

### Strict Priority Queuing

With strict priority queuing, a higher priority queue must be empty before a lower priority queue can transmit any packets. Strict priority queuing is depicted in Figure 40-2. Inbound packets enter on the upper left and proceed to the appropriate queue, based upon the TxQ configuration in the CoS. Outbound packets exit the queues on the lower right. At this time only queue 3 packets are forwarded. This will be true until queue 3 is completely empty. Queue 2 packets will then be forwarded. Queue 1 packets will only forward if both queue 2 and queue 3 are empty. Queue 0 packets will only forward if all other queues are empty. Strict priority queuing assures that the highest priority queue with any packets in it will get 100 percent of the bandwidth available. This is particularly useful for one or more priority levels with low bandwidth and low tolerance for delay. The problem with strict priority queuing is that should the higher level queues never fully empty, lower level queues can be starved of bandwidth.

**Figure 40-2 Strict Priority Queuing Packet Behavior**



### Low Latency Queuing

A Low Latency Queue (LLQ) is a non-configurable strict priority queue. LLQs are designed to guard against:

- Packet loss
- Delay
- Jitter

LLQ hardware resources can not be configured, but a policy can be configured for a CoS that is mapped to an LLQ. In this way, traffic associated with high value real-time voice or video packets can be mapped to an LLQ. The LLQ priority will determine when mapped traffic will be serviced


relative to other traffic. For example, S-Series queues 0, 9, and 10 are LLQs. If a voice policy is mapped to a CoS with a TxQ reference that is in turn mapped to queue 9, this voice traffic will be serviced as soon as queue 10 is empty and will continue to be serviced ahead of any lower priority queue until there is no traffic left in queue 9.

LLQs are hardware dependent. Use the `show cos port-config txq` command to display LLQs for a given module.

## Weighted Fair Queuing

With weighted fair queuing, queue access to bandwidth is divided up by percentages of the time slices available. For example, if 100 percent is divided into 64 time slices, and each queue is configured for 25 percent, each queue will get 16 time slices, after which the next lowest priority queue will get the next 16, and so on. Should a queue empty before using its current share of time slices, the next priority queue inherits the time slices that remain. [Figure 40-3](#) on page 40-5 depicts how weighted fair queuing works. Inbound packets enter on the upper left of the box and proceed to the appropriate priority queue. Outbound packets exit the queues on the lower right. Queue 3 has access to its percentage of time slices so long as there are packets in the queue. Then queue 2 has access to its percentage of time slices, and so on round robin. Weighted fair queuing assures that each queue will get at least the configured percentage of bandwidth time slices. The value of weighted fair queuing is in its assurance that no queue is starved for bandwidth. The downside of weighted fair queuing is that packets in a high priority queue, with low tolerance for delay, will wait until all other queues have used the time slices available to them before forwarding. So weighted fair queuing would not be appropriate for applications with high sensitivity to delay or jitter, such as VoIP.

**Figure 40-3 Weighted Fair Queuing Packet Behavior**






## Hybrid Queuing

Hybrid queuing combines the properties of both strict priority and weighted fair queuing. [Figure 40-4](#) on page 40-6, depicts hybrid queuing. The configuration is for strict priority queuing on queue 3 and weighted fair queuing for the remaining queues, with queue 2 receiving 50 percent of the remaining time slices, and the other queues receiving 25 percent each. The benefit of hybrid queuing is that queues configured as strict priority will receive all the bandwidth that is available in the order of their priority until empty. Remaining bandwidth will be used by the weighted fair queues based upon the time slice percentages configured. The down side remains that anytime strict priority queuing is used, should the strict priority queues never fully empty, remaining queues will be starved of bandwidth.

**Figure 40-4 Hybrid Queuing Packet Behavior**



## Enhanced Transmission Selection

Enhanced Transmission Selection (ETS) queuing provides for configuring two or more traffic class queues (transmit queue (TxQ)) to be allocated for bandwidth that will not be serviced until all non-ETS queues are empty. See [“Enhanced Transmission Selection Configuration”](#) on page 16-2 for ETS feature details.


## Rate Limiting

Rate limiting is used to control the rate of traffic entering (inbound) and/or leaving (outbound) a switch per CoS. Rate limiting allows for the throttling of traffic flows that consume available bandwidth, in the process providing room for other flows. Rate limiting guarantees the availability of bandwidth for other traffic by preventing the rate limited traffic from consuming

more than the assigned amount of a network's resources. Rate limiting accomplishes this by setting a cap on the bandwidth utilization of specific types of both inbound and outbound traffic.

Figure 40-5 on page 40-7 illustrates how bursty traffic is clipped above the assigned threshold with rate limiting applied.

**Figure 40-5 Rate Limiting Clipping Behavior**




## Rate Shaping

Rate Shaping throttles the rate at which a port transmits (outbound) queued packets. Rate Shaping buffers packets received above the configured rate on a per CoS basis, rather than dropping them. Only when buffer capacity is exceeded are packets dropped. Rate shaping may be configured for a CoS on a port, for an 802.1p priority on a port, or for all Classes of Service on a port.

Figure 40-6 illustrates how bursty traffic is smoothed out when it bursts above the assigned threshold with rate shaping applied.

**Figure 40-6 Rate Shaping Smoothing Behavior**



Rate shaping retains excess packets in a queue and then schedules these packets for later transmission over time. Therefore, the packet output rate is smoothed and bursts in transmission are not propagated as seen with rate limiting.

Rate shaping can be implemented for multiple reasons, such as controlling bandwidth, to offer differing levels of service, or to avoid traffic congestion on other links in the network by removing

the burstiness property of traffic that can lead to discarded packets. Rate shaping is important for real-time traffic, where packet loss is extremely detrimental to these applications. Instead of discarding traffic imposed by rate limiting, delays are induced into its transmission by retaining the data for future transmission. However, the delays must also be bounded to the degree that the traffic is sensitive to delays.

## Understanding QoS Configuration on the 7100-Series

This section discusses the six components for configuring QoS and displaying QoS status on an 7100-Series switch router:

**CoS Port-Type:** Based upon the transmit queue (TxQ), Inbound Rate Limiting (IRL), and flood control resource capabilities of the ports in your system. Knowledge of these capabilities is important when configuring queue behaviors. Port group membership and the port resources available are determined by port type.

**CoS Port Groups:** Provide for the grouping of ports by the same class of service features and port type.

**CoS Port Resource Table:** Enables the association of rate limiter and rate shaper values to a port.

**CoS Reference Mapping Table:** Maps your defined TxQ and IRL index references, used by the CoS settings table, to the physical queue and rate limiter settings created in the port-resource table.

**CoS Settings Table:** Used for CoS parameter assignment and contains the current settings for each class of service feature. Each class of service entry consists of an entry index, an 802.1p priority, an optional ToS value, a transmit queue reference, an IRL reference, and a flood control reference. On the 7100-Series flood control is not configurable. Flood control defaults to enabled for all CoS indexes and is applied when a resource is defined.

**CoS State:** A global setting that must be enabled for a configured CoS to affect port behavior. When enabled, CoS state associated with a port supersedes current default or modified port-level controls for priority queue mapping, port rate limiting, and transmit queue. When disabled the port settings apply.



**Note:** It is recommended that you use Extreme Networks NetSight Policy Manager as an alternative to CLI for configuring policy-based CoS on Extreme Networks Series devices.

A policy discussion is outside the scope of this document and will be limited to the relevant configuration example commands. See [Chapter 22, Policy Configuration](#) for a detailed policy discussion.

Numerous QoS values are associated with each other through reference. With the exception of 802.1p priority and ToS, CoS values are first mapped to a port group, which associates a CoS configuration with a port type. A port group has the following CoS parameters associated with it:

- Physical port(s)
- Strict priority or weighted fair queuing behavior
- Rate-limit setting(s)
- Rate-shaping setting(s)
- A port queue
- A port reference

Understanding how these parameters are first mapped to the port group and then to a TxQ or IRL reference is the key to understanding QoS configuration. Where appropriate, the task column in [Procedure 40-1](#) on page 40-17 identifies these mapping relationships.

See “[Determining CoS Port-Type](#)” on page 40-9 and “[Configuring CoS Port Groups](#)” on page 40-10 for a port group discussion.

## Determining CoS Port-Type

Based on physical capability, all physical ports belong to one of two port-types. The importance of this port-type distinction lies in the resources available for transmit queue, inbound rate limiting, outbound rate limiting, and flood control CoS features. The nomenclature distinguishes the type as port type 0.

### TxQ

Port type 0 supports eight transmit queues. Use the **show cos port-type txq** to display all the system's ports currently associated to each type.

The following example displays default values for the **show cos port-type txq** command output:

```
7100 Chassis(rw) -> show cos port-type txq
```

```
Number of resources:          Supported rate types:
txq = transmit queue(s)      perc = percentage
irl = inbound rate limiter(s) pps = packets per second
orl = outbound rate limiter(s) Kbps = kilobits per second
fld = flood rate limiter(s)  Mbps = megabits per second
                              Gbps = gigabits per second
                              Tbps = terabits per second
```

Index	Port type description	Number of queues / slices	Supported rate type	Eligible ports	Unselected ports
0	TOR-Series 8Q	100/8	perc Kbps Mbps Gbps	tg.1.1-48; fg.1.1-4	tg.1.1-48; fg.1.1-4

### IRL

Type 0 supports 16 Inbound Rate Limiters. Use the **show cos port-type irl** command to display the port types and their associated ports.

The following example displays default values for the **show cos port-type irl** command output:

```
7100 Chassis(rw) -> show cos port-type irl
```

```
Number of resources:          Supported rate types:
txq = transmit queue(s)      perc = percentage
irl = inbound rate limiter(s) pps = packets per second
orl = outbound rate limiter(s) Kbps = kilobits per second
fld = flood rate limiter(s)  Mbps = megabits per second
                              Gbps = gigabits per second
                              Tbps = terabits per second
```

Index	Port type description	Number of limiters	Supported rate type	Eligible ports	Unselected ports
0	TOR-Series 16 IRL	16 irl	perc pps	tg.1.1-48; fg.1.1-4	tg.1.1-48; fg.1.1-4

Kbps  
Mbps  
Gbps

## Flood Control

Flood Control is supported on port-type 0. Three reference limiters are supported. Use the **show cos port-type flood-ctrl** command to display the port types and their associated ports.

The following example displays default values for the **show cos port-type flood-ctrl** command output:

7100 Chassis(rw) ->**show cos port-type flood-ctrl**

```
Number of resources:          Supported rate types:
txq = transmit queue(s)      perc = percentage
irl = inbound rate limiter(s) pps = packets per second
orl = outbound rate limiter(s) Kbps = kilobits per second
fld = flood rate limiter(s)  Mbps = megabits per second
                              Gbps = gigabits per second
                              Tbps = terabits per second
```

Index	Port type description	Number of limiters	Supported rate type	Eligible ports	Unselected ports
0	TOR-Series Flood Ctrl	3 fld	pps	tg.1.1-48; fg.1.1-4	tg.1.1-48; fg.1.1-4

## Configuring CoS Port Groups

CoS port groups provide for grouping ports by CoS feature configuration and port type. Ports are required to be configured by groups: this feature provides a meaningful way of identifying ports by similar functionality and port type.

Groups consist of a group number and port type and are numbered as such, *port-group.port-type*. For example: port group 0, port type 0 would be numbered port group 0.0. Three default port groups exist for TxQ, IRL, and flood control CoS features and are identified as port group 0 and port type 0 and are indexed as 0.0 for each feature. These default port groups cannot be removed and all physical ports in the system are assigned to one of the three port groups for each feature (remember group assignment is determined by port type).

Additional port groups, up to eleven total, may be created. Ports assigned to a new port group cannot belong to another non-default port group entry and must be comprised of the same port type as defined by the port group you are associating it with. The creation of additional port groups could be used to combine similar ports by their function for flexibility. For instance, ports associated to users can be added to a port group called Users and ports associated to uplink ports can be added to a port group called Uplink. Using these port groups, a class of service unique to each group can assign different rate limits to each port group. User ports can be assigned a rate limit configured in one CoS, while Uplink ports can be assigned a different rate limit configured in another CoS. A maximum of 8 port groups per CoS transmit queue and/or rate-limiter function are supported.

### Port-Groups: TxQ Configuration

TxQ Port-Groups contain user settings for specific types of ports and their matching transmit queue settings. Port group 0.0 exists by default. New port groups can be configured with a name and ports can be added according to device port-type. Transmit queue behavior can also be

configured per port group; default port-groups are configured in strict priority queuing mode. Additional port groups also default to strict priority queuing mode, though each TxQ port group can be configured for weighted-fair queuing if desired.

The **show cos port-config txq** command displays all configured TxQ port-groups. Group name and type are displayed as well as ports associated with the port group. For **show cos port-config txq** output, arbiter mode (TxQ mode) is displayed along with a picture of the supported queues and the number of slices allotted to the group. Queuing is also displayed by percentage.

The following example displays default values for the **show cos port-config txq** command output:

```
7100 Chassis(rw)->show cos port-config txq

* Percentage/queue (if any) are approximations based on
  [(slices/queue) / total number of slices]

Transmit Queue Port Configuration Entries
-----
Port Group Name   :TOR-Series 8Q
Port Group       :0
Port Type        :0
Assigned Ports   :tg.1.1-48;fg.1.1-4
Arbiter Mode     :Strict
Slices/queue     :Q [ 0]:  0 Q [ 1]:  0 Q [ 2]:  0 Q [ 3]:  0
                  :Q [ 4]:  0 Q [ 5]:  0 Q [ 6]:  0 Q [ 7]: 100
Percentage/queue :Q [ 0]:  0% Q [ 1]:  0% Q [ 2]:  0% Q [ 3]:  0%
                  :Q [ 4]:  0% Q [ 5]:  0% Q [ 6]:  0% Q [ 7]: 100%
-----
```

Additional port groups can be created using the **set cos port-config txq** command. Name and associated ports can be configured, as well as TxQ settings. You need to:

- Identify the port-group for configuration
- Optionally, specify port-group Name, associated ports, and arb-percentage or arb-slices

## Port-Groups: IRL Configuration

IRL port-groups contain user settings for specific types of ports and their matching inbound rate limiting configurations. Port groups 0.0 through 0.1 exist by default. Each new group can be configured with a name and ports added to each group according to device port-type. Use the **show cos port-config irl** command to display each IRL port-group configured by group and type, with group name and associated ports.

The following example displays default values for the **show cos port-config irl** command output:

```
7100 Chassis(rw)->show cos port-config irl

Inbound Rate Limiting Port Configuration Entries
-----
Port Group Name   :TOR-Series 16 IRL
Port Group       :0
Port Type        :0
Assigned Ports   :tg.1.1-48;fg.1.1-4
-----
```

Additional port groups can be created using the **set cos port-config irl** command. Port group name and associated ports can be configured. You need to:

- Identify the port-group for configuration
- Optionally, specify port-group Name and associated ports

## Port-Groups: Flood Control Configuration

CoS-based flood control prevents configured ports from being disrupted by a traffic storm by rate limiting specific types of packets through those ports. When flood control is enabled on a port, incoming traffic is monitored over one second intervals. During an interval, the incoming traffic rate for each configured traffic type (unicast, broadcast, or multicast) is compared with the configured traffic flood control rate, specified in packets per second. If, during a one second interval, the incoming traffic of a configured type reaches the traffic flood control rate configured on the port, CoS-based flood control drops the traffic until the interval ends. Packets are then allowed to flow again until the limit is again reached.

Flood control port-groups contain user settings for specific types of ports and their matching flood limiting configurations. Port group 0.0 exists by default. Each new group can be configured with a name and ports added to each group according to device port-type. Use the **show cos port-config flood-ctrl** command to display each flood control port-group configured by group and type, with group name and associated ports.

The following example displays default values for the **show cos port-config flood-ctrl** command output:

```
7100 Chassis(rw)->show cos port-config flood-ctrl
```

```
Flood Rate Limiting Port Configuration Entries
```

```
-----
Port Group Name   :TOR-Series Flood Ctrl
Port Group       :0
Port Type        :0
Assigned Ports   :tg.1.1-48;fg.1.1-4
-----
```

Additional port groups can be created using the **set cos port-config flood-ctrl** command. Port group name and associated ports can be configured. You need to:

- Identify the port-group for configuration
- Optionally, specify port-group Name and associated ports

## Configuring CoS Port-Resource

Physical rate limiters and rate shapers are configured in CoS port resources. Resources map directly to the number of queues and rate limiters supported by each port-type. Remember, group 0.0 supports 16 TxQ resources and 32 IRL resources. Resources exist for each port group and are indexed as *port-group.port-type resource-index*. Port-resources initially default to none, as rate limiting and shaping is not required.

### CoS TxQ Port-Resource (Outbound Rate Shapers)

Rate shaping throttles the rate at which queues transmit packets. See “[Rate Shaping](#)” on page 40-7 for a general discussion of rate shaping. Rate shaping is TCP friendly; it buffers packets that are above the rate rather than drop them. CoS rate shaping allows you to configure rate shapers based on a unit rate (kilobits/second, megabits/second, gigabits/second), or a percentage of the port’s line speed.

The **show cos port-resource txq** command displays resources for each port group created along with the resource index (physical queue). By default, no resources are configured for TxQ port-resources. Rates displayed as none indicate no resources exist. The default Rate Shaping algorithm is tail-drop and is not configurable.

The following example displays default values for the **show cos port-resource txq** command output:

```
7100 Chassis(rw)->show cos port-resource txq
```

'?' after the rate value indicates an invalid rate value

Group	Index	Resource	Type	Unit	Rate	Algorithm
0.0	0	txq	perc	none		tail-drop
0.0	1	txq	perc	none		tail-drop
0.0	2	txq	perc	none		tail-drop
0.0	3	txq	perc	none		tail-drop
0.0	4	txq	perc	none		tail-drop
0.0	5	txq	perc	none		tail-drop
0.0	6	txq	perc	none		tail-drop
0.0	7	txq	perc	none		tail-drop

'?' after the rate value indicates an invalid rate value

The **set cos port-resource txq** command is used for creating outbound rate shapers. You need to:

- Identify the port group for configuration
- Identify the queue resource ID, along with unit and rate desired for that queue

## CoS IRL Port-Resource (Inbound Rate Limiter)

Unlike rate shaping, inbound rate limiting or rate policing simply drops or clips traffic inbound if a configured rate is exceeded. See [“Rate Limiting”](#) on page 40-6 for a general discussion of rate limiting. CoS inbound rate limiting allows you to configure rate limits based on a unit rate (kilobits/second, megabits/second, gigabits/second), or percentage of the port’s line speed.

The **show cos port-resource irl** command displays resources for each port group created along with the index, as described above. By default, no resources are configured for IRL port-resources. Rates displayed as none indicate no resources exist. The default Rate Limiting algorithm is tail-drop.

The following example displays default values for the **show cos port-resource irl** command output:

```
7100 Chassis(rw)->show cos port-resource irl
```

'?' after the rate value indicates an invalid rate value

Group	Index	Resource	Type	Unit	Rate	Rate Limit	Type	Action
0.0	0	irl	perc	none		drop		none
0.0	1	irl	perc	none		drop		none
0.0	2	irl	perc	none		drop		none
0.0	3	irl	perc	none		drop		none
0.0	4	irl	perc	none		drop		none
0.0	5	irl	perc	none		drop		none
0.0	6	irl	perc	none		drop		none
0.0	7	irl	perc	none		drop		none
0.0	8	irl	perc	none		drop		none

.



.

The **set cos port-resource irl** command is used for creating inbound rate limiters. You need to:

- Identify the port group for configuration
- Identify the limiter resource ID, along with desired unit, rate, and actions

### CoS Flood Control Port-Resource (Flood Limiter)

Flood control limiting prevents configured ports from being disrupted by a traffic storm by rate limiting configured traffic types such as multicast or broadcast through those ports. CoS flood limiting allows you to configure traffic type limiting based on a unit rate (kilobits/second, megabits/second, gigabits/second), or percentage of the port’s line speed.

The **show cos port-resource flood-ctrl** command displays resources for each port group created along with the index, as described above. By default, no traffic type is configured for flood control port-resources. Rates displayed as none indicate no resources exist. The default rate limiting algorithm is tail-drop.

The following example displays default values for the **show cos port-resource flood-ctrl** command output:

```
7100 Chassis(rw)->show cos port-resource flood-ctrl
```

'?' after the rate value indicates an invalid rate value

Group	Index	Resource	Type	Unit	Rate	Rate Limit	Type	Action
0.0	1	f1d	pps	none				none
0.0	2	f1d	pps	none				none
0.0	3	f1d	pps	none				none

Configure a CoS flood control resource entry, by mapping a port group with a traffic type such as multicast or broadcast. This index is used by the rate-limit option when setting a flood control cos reference

The **set cos port-resource flood-ctrl** command is used for configuring a CoS flood control resource entry.

### Configuring CoS Reference Mapping

The CoS Reference Table maps the TxQ and IRL references, defined by you and configured in the CoS Settings Table, to physical queues and rate limiters created in the port-resource table. A CoS reference table exists for each port group. The CoS reference table indexes can be thought of as virtual queues or rate limiters. The table accounts for the maximum number of queues and rate limiters supported by the device. The virtual queues and limiters map to the physical queues and rate limiters. The TxQ reference table is populated by default, because queues are required for all forwarding. The TxQ reference maps each reference value to a physical queue. The IRL Reference Table is not configured by default, because inbound rate limiting is optional.

### CoS TxQ Reference Mapping

The CoS TxQ reference table uses 16 indexes or virtual queues, and maps each to a physical queue or resource. A TxQ reference table exists for each port group configured and is indexed similarly to port-resources, as *port-group.port-type reference*. For port-types with 16 queues, the 16-txq reference indexes map directly to the 16 physical queues. For port-types with 4 queues, the 16-txq reference indexes map:

- virtual queues 12-15 to physical queue 3

- virtual queues 8-11 map to physical queue 2
- virtual queues 4-7 map to physical queue 1
- virtual queues 0-3 map to physical queue 0

The TxQ reference table can be displayed using the **show cos reference txq** command and displays port-group, reference index, and physical queue.

Although the TxQ reference table is populated by default, the Queue-to-Reference mapping can be configured using the **set cos reference txq** command. You need to:

- Identify the port group for configuration
- Identify the transmit queue reference, along with the associated queue

## CoS IRL Reference Mapping Table

The CoS IRL reference table uses 32 indexes or virtual rate limiters, and maps each virtual limiter to a physical limiter or resource. An IRL reference table exists for each port group configured, and is indexed similarly to port-resources, as *port-group.port-type reference*. Because it is an optional configuration, IRL references are not populated with limiters (resources), but can be configured by you. The IRL reference table can be displayed using the **show cos reference irl** command.

The following example displays default values for the **show cos reference irl** command output:

```
7100 Chassis(rw) ->show cos reference irl
```

Group	Index	Reference	Type	Rate Limiter
0.0	0		irl	none
0.0	1		irl	none
0.0	2		irl	none
0.0	3		irl	none
0.0	4		irl	none
0.0	5		irl	none
0.0	6		irl	none
0.0	7		irl	none
0.0	8		irl	none
.				
.				
.				

Physical-Limiter to reference mapping can be configured using the **set cos reference irl** command. The other references not configured are indicated by rate limiter “none”. To configure a physical limiter to reference mapping, you need to:

- Identify the port group for configuration
- Identify the rate-limit reference

## Configuring the CoS Index

The CoS settings table assigns a priority, a ToS value, TxQ reference table and an IRL reference to a CoS entry as follows:

**CoS Index** - Indexes are unique IDs for each CoS settings table entry. CoS indexes 0–7 are created by default and mapped directly to an 802.1p priority values 0–7 for backwards compatibility. These entries cannot be removed and the 802.1p value cannot be changed. When CoS is enabled using the **set cos state enable** command, indexes are assigned. Entries 0–255 are configurable for a total of 256 CoS entries.

**Priority:** For each new CoS index created, you have the option to assign an 802.1p priority value 0-7 for the class of service. CoS indexes 0-7 map directly to 802.1p priorities and cannot be changed as they exist for backward compatibility. All other CoS index entries can have a priority value set between 0 and 7.

**ToS:** The IP header Type of Service field is an 8-bit field also referred to as the DiffServ Code Point (DSCP) field. This optional value can be set per class of service to a value between 0–255. When a frame is assigned to a class of service for which this value is configured, the ToS field of the incoming IP packet will be overwritten to values defined by you. This ToS rewrite option also allows masking. The ToS can selectively mask (change) certain bits of the field, without changing others. For instance, masking the ToS could be used to modify the ToS precedence without modifying the DTR/ECN bits. The mask specified contains the bits to be changed. CLI input can be in decimal or hex value, and a mask is not required. If the mask is not specified in the ToS input, all bits will be overwritten. ToS can be set for CoS indexes 0-7.

**TxQ Reference:** Because all traffic requires association to a transmit queue, the CoS TxQ reference field will always be populated when a new CoS index is created. If a TxQ reference value is not chosen, TxQ reference 0 will be assigned. The reference does not indicate the actual transmit queue to be assigned by CoS; it points to the CoS TxQ reference mapping table index entry. It may be thought of as the virtual queue that is associated to a physical queue defined by the TxQ reference mapping table. TxQ reference mapping table defines 16 TxQ references, therefore CLI input for TxQ reference in the CoS Settings Table is 0-15. See “[CoS TxQ Reference Mapping](#)” on page 40-14 for a TxQ reference configuration discussion.

**IRL Reference:** The CoS IRL reference field is optional, as rate limits are not required. Like the TxQ reference field, the IRL reference does not assign an inbound rate limit but points to the CoS IRL Reference Mapping Table. This reference may also be thought of as the virtual rate limiter that will assign the physical rate limiter defined by the IRL Reference Mapping Table. The IRL Reference Mapping Table defines 32 IRL references, therefore input for IRL reference in the CoS Settings Table is 0-31. See “[CoS IRL Reference Mapping Table](#)” on page 40-15 for an IRL reference configuration discussion.

New CoS Indexes can be created using the **set cos settings** command. ToS, 802.1p priority, TxQ reference, and IRL Reference can be configured for each CoS Index. You need to:

- Enter a CoS Index value from 0–255
- Specify 802.1p priority (Index entries 8–255 only), tos-value, txq-reference and irl-reference

Use the **set cos settings** command to create or modify an already existing CoS index.

Use the **show cos settings** command to display current CoS indexes.

The following example displays default values for the **show cos settings** command output:

```
7100 Chassis(rw) ->show cos settings
```

```
* Means attribute has not been configured
```

CoS Index	Priority	ToS	TxQ	IRL	Flood-Ctrl
0	0	*	0	*	Enabled
1	1	*	1	*	Enabled
2	2	*	2	*	Enabled
3	3	*	3	*	Enabled
4	4	*	4	*	Enabled
5	5	*	5	*	Enabled
6	6	*	6	*	Enabled
7	7	*	7	*	Enabled

## Enabling CoS State

CoS state is a global setting that must be enabled for CoS configurations to be applied to a port. When CoS state is enabled, controls configured for CoS supersede port level controls for priority queue mapping, IRL, and TxQ. These port level settings can be configured independent of CoS state, but will have no affect while CoS is enabled. Disabling CoS results in the restoration of current port level settings.

Use the **set cos state enable** command to enable CoS state globally for this system.

Use the **set cos state disable** command to disable CoS state globally for this system.

Use the **show cos state** command to display the current status of CoS state.

## The QoS CLI Command Flow

[Procedure 40-1](#) provides a CLI flow summary of each step in the configuration flow along with the show commands to verify the configuration.

### Procedure 40-1 Class of Service CLI Configuration Command Summary

Step	Task	Command(s)
1.	Inspect the TxQs, IRL, and flow control support for the installed ports. This information is used to determine the module port type for port group.	<b>show cos port-type txq</b> <b>show cos port-type irl</b> <b>show cos port-type flood-ctrl</b>
2.	Set the CoS transmit queue port group configuration by mapping a physical port list to a port group for purposes of TxQ configuration. Optionally associate a name and the configuration of a TxQ weighted fair queue behavior configuration. Verify the new configuration.	<b>set cos port-config txq</b> <i>group-type-index</i> [ <i>name name</i> ] [ <i>ports port-list</i> ] [ <i>append</i> ]   [ <i>clear</i> ] [ <i>arb-slice slice-list</i> ] [ <i>arb-percentage percentage-list</i> ] [ <i>enhanced-groups group-id</i> ] [ <i>enhanced-percentage bandwidth</i> ] <b>show cos port-config txq</b> <i>port_group.port_type</i>
3.	Set the CoS inbound rate-limit port group configuration by mapping a physical port list to a port group for purposes of IRL configuration, optionally allowing the association of a name for this configuration. Verify the new configuration.	<b>set cos port-config irl</b> <i>port_group.port_type</i> <i>name name</i> <i>ports ports_list</i> <b>show cos port-config irl</b>
4.	Set the CoS flood control limit port group configuration by mapping a physical port list to a port group for purposes of flood control configuration, optionally allowing the association of a name for this configuration. Verify the new configuration.	<b>set cos port-config flood-ctrl</b> <i>port_group.port_type</i> <i>name name</i> <i>ports ports_list</i> <b>show cos port-config flood-ctrl</b>
5.	Configure a Class of Service transmit queue port resource entry, by mapping a port group with a transmit queue and applying a TxQ rate shaping value to the mapping. Verify configuration changes.	<b>set cos port-resource txq</b> <i>port_group.port_type</i> <i>tx_queue</i> <b>unit</b> <i>unit</i> <b>rate</b> <i>rate</i> <b>show cos port-resource txq</b> <i>port_group.port_type</i>
6.	Configure a CoS inbound rate limiting index entry, by mapping a port group with a rate-limit value. This index is used by the rate-limit option when setting an IRL cos reference.	<b>set cos port-resource irl</b> <i>port_group.port_type</i> <i>index</i> <b>unit</b> <i>unit</i> <b>rate</b> <i>rate</i> <b>show cos port-resource irl</b> <i>port_group.port_type</i>

**Procedure 40-1 Class of Service CLI Configuration Command Summary (continued)**

Step	Task	Command(s)
7.	Set a CoS inbound rate limiting reference configuration, by mapping a port group with a rate limiter resource ID and associating the mapping with an IRL reference. Verify the new CoS reference configuration.	<b>set cos reference</b> <i>irl port_group.port_type reference rate-limit IRLreference</i> <b>show cos reference</b> <i>irl port_group.port_type</i>
8.	Modify a currently configured CoS or create a new CoS. Verify the new CoS configuration. All TxQ to port group mappings are associated with the transmit queue reference. All IRL to port group mappings are associated with the inbound rate limiter reference.	<b>set cos settings</b> <i>cos-list [priority priority] [tos-value tos-value] [txq-reference txq-reference] [irl-reference irl-reference]</i> <b>show cos settings</b>
9.	Enable CoS state for the system. Verify the new CoS state.	<b>set cos state enable</b> <b>show cos state</b>

## Terms and Definitions

Table 40-1 lists terms and definitions used in this Quality of Service configuration discussion.

**Table 40-1 Quality of Service Configuration Terms and Definitions**

Term	Definition
Class of Service (CoS)	The grouping of priority and forwarding behaviors that collectively determine packet bandwidth behavior as it transits the link, including: 802.1p, IP ToS rewrite, priority Transmit Queue (TxQ), Inbound and/or outbound Rate Limiter (IRL) and outbound rate shaper.
DSCP	Differentiated Services Code Point. The lower 6 bits of the ToS field defined by RFC 2474.
Flows	In a QoS context, a sequence of IP packets that share a common class of service and forwarding treatment as they transit the interface.
Forwarding Treatment	Queue behavior during the packet egress stage (strict priority, weighted fair, hybrid).
Jitter	The change in a flow's packet spacing on the link due to the bursty and congestive nature of the IP network. This irregular spacing - jitter - can severely degrade the quality of voice calls or multimedia presentations.
Port Group	The grouping of ports based upon the same CoS features and port type.
Port Type	The differentiation of ports based upon TxQ, IRL, and flood control resource capabilities.
Priority	The preference of one packet (classification) or queue (packet forwarding) over another.
Quality of Service (QoS)	A bandwidth management mechanism able to preferentially treat packets based upon packet classification and forwarding treatment.
Rate Limiting	The bounding of bandwidth used by a QoS packet flow such that excess packets are dropped/clipped.
Rate Shaping	The rescheduling of bursty traffic while in the queue based upon packet buffering such that traffic beyond the configured bandwidth threshold is delayed until bandwidth usage falls below the configured threshold.

**Table 40-1 Quality of Service Configuration Terms and Definitions (continued)**

Term	Definition
Type of Service (ToS)	An 8-bit field defined by RFC 1349 used for the prioritization of packets within a QoS context.



## Anti-Spoofing Configuration

This chapter describes the anti-spoofing features and how to configure them.

For information about...	Refer to page...
<a href="#">Anti-Spoofing Feature Overview</a>	41-1
<a href="#">Implementing Anti-Spoofing in Your Network</a>	41-4
<a href="#">Anti-Spoofing Configuration</a>	41-5

### Anti-Spoofing Feature Overview

Attacks on IP networks can easily be performed using readily available tools found on the internet today. Malicious users can spoof DHCP server response packets, allowing them to give false information to a user for such fields as the default gateway or domain name resolution servers. Man in the middle attacks can take advantage of ARP, allowing a hacker to redirect user traffic through his own device to and from the default gateway. The hacker can then spy on the private information being sent from the user, without either the user or gateway knowing. A malicious user can spoof an innocent user's IP address, allowing the malicious user to bypass other possible security features of a network that are based on a user's subnet.

The Extreme Networks anti-spoofing solution provides a flexible and secure approach to IP spoofing detection and prevention. To mitigate the effects of these types of attacks on a network, a source MAC to source IP address binding table is created. The three basic tools used to detect source IP to source MAC address associations, based on the entries in the binding table, and take action on violations are:

- DHCP snooping,
- Dynamic ARP inspection (DAI), and
- IP source guard.

All three methods can create IP-to-MAC bindings in the binding table, although both DAI and IP source guard can be configured to run in inspection-only mode, limiting the association of IP addresses to MAC addresses to DHCP-snooping. Bindings created as a result of DHCP exchanges with trusted servers (DHCP-snooping) take precedence over bindings created through DAI or IP source guard. Use of all three tools allows bindings to be created for users in a network where DHCP is not in use or where a DHCP exchange has not occurred since the anti-spoofing feature has been enabled.

The actions that may be taken against a violating user include:

- Logging a message
- Sending a notification
- Putting the user in quarantine, as defined by a policy profile



## DHCP Snooping

DHCP snooping provides the foundation for IP spoofing detection and prevention. DHCP ACK packets received on a trusted port from a DHCP server create a MAC-to-IP binding for the user along with the lease time and expiration. DHCP ACK packets received on any ports that are configured as untrusted should be dropped as configured by policy.

On edge devices, an optional configuration is to verify the SA (source address) MAC address of the client with the client hardware address found in the DHCP payload. Provided that policy is appropriately configured to determine trusted ports for DHCP servers versus DHCP clients in an exclusively DHCP environment, and configured on an edge switch, DHCP snooping is deterministic in binding an IP address to a MAC address. Dynamic ARP inspection and IP source guard can be used to supplement the bindings database and create a secure network.

### DHCP Snooping Port Mode

In a DHCP snooping context, there are three configurable port modes that determine anti-spoofing behavior:

- **Trusted** – When port mode is set to trusted, DHCP server traffic is accepted and used to create bindings in the source MAC address to IP address binding table for the user. Binding verification does not take place on trusted ports.
- **Bypass** – When port mode is set to bypass, snooping of DHCP server traffic does not take place on the port.
- **Untrusted** – When port mode is set to untrusted, the untrusted server counter is incremented when DHCP server traffic is detected on the port. Client traffic on these ports is processed when MAC verification is enabled on these ports.

Bindings created as a result of DHCP exchanges on trusted ports using DHCP snooping take precedence over bindings created through dynamic ARP inspection or IP source guard.

### DHCP Snooping MAC Verification

The DHCP client packet contains an L2 source MAC address and an L3 client hardware address. When DHCP snooping MAC verification is enabled, DHCP snooping verifies that the source MAC address and the client hardware address match in DHCP client packets that transit untrusted ports. If the addresses do not match, the packet is dropped.

DHCP MAC verification is a network edge feature that should be enabled on ports transited by client packets from the intended client. For DHCP snooping MAC verification to be operational:

- DHCP snooping must be enabled, globally and on the port
- The port mode must be set to untrusted

## Dynamic ARP Inspection (DAI)

Dynamic ARP inspection uses the MAC-to-IP binding database to ensure that ARP packets have the proper MAC-to-IP binding. When an ARP packet enters the switch, the source MAC and IP addresses are compared to the entry in the table. If the packet data conflicts with the binding in the table, the IP change is counted and logged, and any configured actions are taken against the user.

DAI can also be configured to populate the MAC-to-IP binding table. Successfully limiting ARPs to the bound addresses in the table prevents a malicious user from inserting himself in between the end user and a gateway and poisoning network devices' ARP caches or succeeding in MITM (man in the middle) attacks.

## IP Source Guard

IP source guard is another means to restrict IP traffic and take configured actions against violating users. IP traffic on a port is inspected to ensure that a user's MAC and IP addresses are found in the binding table created by DHCP snooping. Changes to a user's IP address are counted and action is taken, as configured.

Like DAI, the anti-spoofing feature can be configured so that IP source guard is also able to add entries to the MAC-to-IP binding database dynamically, based upon IP traffic traversing the switch. This is particularly beneficial in an environment not limited to edge devices or one in which DHCP is not the sole proprietor of network IP addresses.

## Duplicate IP Address Detection

In addition to the anti-spoofing tools described above, the anti-spoofing feature can also be configured to log, through SYSLOG and SNMP traps, duplicate IP addresses when they are bound to different MAC addresses. This situation is usually due to a misconfiguration in the network and is generally not indicative of an attack, but can be a worthwhile event to record, as administrative action may be needed to reconcile the condition. These duplicate IP addresses are only detected upon a user's binding change, and do not apply to duplicate IP addresses over ports for the same MAC address (for example, if a single user moves from one port to another).

## Populating the MAC-to-IP Binding Table

The anti-spoofing MAC-to-IP binding table can be populated through DHCP snooping, dynamic ARP inspection, and IP source guard. Regardless of which of these three methods are adding entries to that table, an entry cannot be added if there is not already an entry for the user's MAC address in the multiauth session table. (Refer to the chapter entitled "Authentication Configuration" in this book for more information about the session table.)

### Bindings Created by DHCP Snooping

DHCP snooping watches DHCP exchanges to create a MAC to IP binding for a client. A basic DHCP client/server exchange is as follows:

1. client -> server: DISCOVER
2. server -> client: OFFER
3. client -> server: REQUEST
4. server -> client: ACKNOWLEDGE

It is the acknowledgement from the server that creates the binding, and the server message is considered authoritative. (No other security measures other than those described here are used to ensure that the server is legitimately responding to a client request.) The ACK message includes the client hardware address and the client's confirmed IP address. It is the client hardware address (not the MAC destination address) that is used in determining if there is already an entry in the multiauth session table for the user, to which the IP address will be bound. If there is no entry in the session table for the client, a message will be logged.

Only DHCP server ACK messages received on trusted ports will populate MAC-IP address bindings. On untrusted ports, any DHCP server packets are recorded (that is, the counter is incremented), but they are allowed to be further processed. If policy is properly configured, the packets will be dropped or the port will be shut down, as configured. These server messages are not used to populate MAC to IP bindings. Bypass ports ignore all DHCP server packets for purposes of populating the binding database.

DHCP server messages are limited to trusted ports, so the bindings that are created by them are not intended to be recorded as violations. In the case that a server sends a client a new binding (with a different IP) before the current binding's lease has expired, the event will trigger a SYSLOG message, but will not increment the violation counter.

If neither DAI nor IP source guard are configured to populate the bindings table (disabled or inspection-only), DHCP server ACK message are required to create the IP binding for a user. In this configuration, the switch will drop any DHCP server messages that cannot be processed by the soft path. The expected result of this would be for the DHCP client to re-initiate a request to the server and thus give the switch another opportunity to add the entry to the binding table. When either DAI or IP source guard is configured to populate the bindings table this is not necessary (the switch can forward the ACK), as the user's binding is able to be populated by other traffic sourced from the user, so the packet would not be discarded.

## Bindings Created by DAI or IP Source Guard

When DAI or IP source guard are enabled, the other traffic being inspected (ARP or IP) can also populate the IP address bindings table. With ARP inspection, the sender MAC and IP and target MAC and IP from the ARP payload are used to populate the bindings, as provided by the ARP request or reply. With IP inspection, the source MAC address and IP address are used in creating these bindings.

If a binding already exists for a user due to DHCP, and the lease time has not expired, the DHCP binding takes precedence and a violation is recorded, but the binding does not change. If there is an entry for the user in the multiauth session table and DHCP snooping has not provided a MAC to IP address binding table entry, the ARP or IP traffic can create the MAC to IP address binding table entry. This form of entry creation allows for the anti-spoofing feature to adapt to environments that are not on the edge or are not able to monitor and process all DHCP exchanges on the network for attached users.

## Expiration of Bindings

IP address bindings will timeout when a lease expires, a DHCP release frame is received, or upon manual clearing of an entry, whichever occurs first. For DHCP-snooping created bindings, after the lease expires, the binding also expires. However, for DAI and IP-inspection, the counter resets after the timeout period, but the binding remains active (restarts the timer).

When you manually set a timeout period, be aware that the lease time defined in the DHCP server scope takes precedence over manually set timeouts.

# Implementing Anti-Spoofing in Your Network

## Using DHCP Snooping Only

On an edge device in an environment where DHCP is exclusively the provider of IP addresses, the switch with DHCP snooping enabled will record all user's DHCP interactions and should have an IP address binding for each connected user.

Untrusted ports do not create bindings from DHCP server packets. Optionally, the client hardware address in the DHCP packet is verified to match the source MAC address of the packet. If it does not, it is dropped. This is a more robust security feature that can be used on the edge of the network where it is expected that the client requests are coming from the client, not a different switch, router, or AP.

No port class actions are taken against users whose IP address assignment changes due to DHCP (where the server responses are on a trusted port), and user counters don't increment. Without DAI or IP source guard configured, anti-spoofing ensures that server packets are only handled

where appropriate, that malicious users do not release or decline DHCP IP address assignments for other users, that DHCP client request packets are coming from the actual client (optional), and that the MAC-IP address binding database is populated. In addition, policy should be configured to drop any unwanted server traffic on untrusted ports.

If IP source guard and DAI are disabled or configured for inspection-only away from the edge of a network, DHCP exchange packets could be missed — for example, link loss at the distribution or core layer would not necessarily cause DHCP renewals from the end users at the edge, thus the binding table would not be repopulated — and users could suffer the consequence of unintended violations (for example, denial of service).

However, there are still benefits for using DHCP snooping without IP source guard or DAI away from the edge of the network. This type of network configuration allows for user accounting (user IP address change counters) and allows for the population of the user IP address binding table from known DHCP servers. The binding table will then allow user leases to run for the configured lease time used on the network before turning on other anti-spoofing features. In this scenario, an administrator should recognize that configuring any actions that limit a user's traffic after a violation could potentially disrupt network traffic for an otherwise legitimate user. Generally, this configuration would not be used away from the edge to quarantine or otherwise limit the users' traffic, as these limitations could be manipulated to cause denial of service attacks against a user.

## Using DAI, IP Source Guard, and Duplicate IP Detection

Once DAI is enabled or set to inspection-only, ARP packet inspection occurs. On those ports, all ARP traffic is intercepted and the MAC and IP address of the ARP is verified against the entry in the MAC to IP address binding table. Actions may be taken against the user if the violation threshold has been crossed for the port, as configured by the port class.

Similarly, if IP source guard is enabled or configured for inspection-only, IP traffic is intercepted and verified against the binding table. Once a connection is created, that traffic won't be inspected again unless the source IP address associated with the MAC address changes. As IP address changes are detected and configured thresholds for that value are crossed, the anti-spoofing feature will take action, depending on the configuration of the class of port with which the user is associated. These actions will be to SYSLOG the event, send an SNMP notification, or perform the quarantine action. The quarantine action is configurable through the policy and multiauth quarantine controls. Extreme Networks **highly recommends that you use quarantine policies to classify the user traffic upon violation hits.**

If the duplicate IP detection feature is enabled, when new MAC to IP bindings are created or current bindings are changed, an IP address lookup is run on the bindings database to verify that the IP is not currently in use. If it is in use, a SYSLOG message and trap are sent.

# Anti-Spoofing Configuration

## Overview

You can enable and disable anti-spoofing on a global and per-port basis. When the feature is globally disabled, no anti-spoofing features are active. Anti-spoofing must be globally enabled before port control values are considered when inspecting traffic. The default value for all anti-spoofing features, global and per port, is disabled.

DHCP snooping is controlled through port enable/disable commands, as well as per port MAC verification enable/disable commands. DAI and IP source guard have individual controls to enable, disable, and enable inspection-only (no binding association) on a per port level. Duplicate IP address detection can be enabled or disabled globally.

Port mode, or type, determines the role traffic traversing the port will take in DHCP snooping. DHCP server messages are only processed (for DHCP snooping purposes) on trusted ports. On untrusted ports, DHCP server messages are counted in the untrusted packet counter (per port). If configured by policy, these message can also be dropped.

On bypass ports, DHCP server messages are ignored (that is, they do not affect the source MAC/source IP binding database, but they are not dropped). Ports are untrusted by default.

## Port Classes

Enabling anti-spoofing on both the global and port level results in snooping frames, but it does not necessarily result in any actions being taken on IP address binding violations. For this, port classes must be defined and ports added to the appropriate port class. Port classes are configured with thresholds and actions, and potentially an action value. Currently, up to 3 port classes can be configured on the switch.

Up to 6 thresholds can be configured per port class, and each threshold can be assigned one of the following actions: sending SYSLOG messages, sending SNMP notifications (traps), and applying the quarantine policy profile. Only the quarantine action can have an action value applied, which is the quarantine profile index. The quarantine profile must be configured independently, and no error checking occurs to ensure the policy profile is present.

Each port can be configured with a single class. If you only have a single anti-spoofing detection type enabled on the port, DHCP snooping for example, the class thresholds and actions can be set for that anti-spoofing detection type. If multiple anti-spoofing types are enabled on a port, DHCP snooping and dynamic ARP inspection for example, the class thresholds and actions must take into account any combination of anti-spoofing events for the configured anti-spoofing types.

If the quarantine action is specified, Extreme Networks highly recommends that you associate a valid quarantine profile with the quarantine action. Refer to the chapter entitled “Policy Configuration” in this book for information about configuring policy profiles and the chapter “Authentication Configuration” for information about using quarantine policies with the quarantine agent.

## Managing the Binding Database

An entry in the source MAC address to source IP address binding table can be deleted by port, source MAC address, or source IP address. Clearing the binding also clears the IP address change count associated with the user. Alternatively, a user’s violation count can be cleared without clearing the current binding.

## Configuration Examples

[Procedure 41-1](#) describes the tasks and commands used to configure anti-spoofing features on the switch. [Table 41-1](#) on page 41-8 describes the tasks and commands used to manage anti-spoofing features. [Table 41-2](#) on page 41-8 describes the commands used to display anti-spoofing information.

Refer to the “Anti-Spoofing Commands” chapter in the 7100-Series *CLI Reference* for details about using these commands.

**Procedure 41-1 Configuring Anti-Spoofing Features**

Step	Task	Command(s)
1.	Create a port class and optionally, configure a name and timeout value. Up to 3 classes can be configured.	<b>set antispoof class</b> <i>class-index</i> { <b>name</b> <i>name</i>   <b>timeout</b> <i>timeout</i> }
2.	Configure thresholds and actions for the class. Up to 6 threshold indexes can be specified per class.	<b>set antispoof class</b> <i>class-index</i> <b>threshold-index</b> <i>thresh-index</i> [ <b>threshold-value</b> <i>thresh-value</i> ] [ <b>quarantine-profile</b> <i>quar-profile</i> ] [ <b>action</b> {[ <b>syslog</b> ] [ <b>trap</b> ]}] [ <b>quarantine</b> ]{}
3.	Enable DHCP snooping on the desired port or ports.	<b>set antispoof dhcp-snooping enable</b> <i>port-string</i>
4.	Configure the ports on which trusted DHCP server traffic will be accepted.  DHCP ACK packets received on these ports will be used to populate the MAC-to-IP address binding table.  All other ports will default to untrusted mode. DHCP packets received on untrusted ports will increment the untrusted server counter.	<b>set antispoof dhcp-snooping port-mode trusted</b> <i>port-string</i>
5.	Optionally, enable DHCP snooping MAC verification on the desired untrusted port or ports.	<b>set antispoof dhcp-snooping mac-verification enable</b> <i>port-string</i>
6.	Optionally, enable dynamic ARP inspection or specify ARP packet inspection only, on the desired port or ports.	<b>set antispoof arp-inspection enable</b> <i>port-string</i> <b>set antispoof arp-inspection inspection-only</b> <i>port-string</i>
7.	Optionally, enable IP source guard or specify IP packet inspection only, on the desired port or ports.	<b>set antispoof ip-inspection enable</b> <i>port-string</i> <b>set antispoof ip-inspection inspection-only</b> <i>port-string</i>
8.	Optionally, configure bypass ports. DHCP server packets received on these ports will be ignored.	<b>set antispoof dhcp-snooping port-mode bypass</b> <i>port-string</i>
9.	Assign port classes to ports.	<b>set antispoof port-class</b> <i>class-index</i> <i>port-string</i>
10.	Globally enable anti-spoofing features on the switch.	<b>set antispoof enable</b>
11.	Optionally, change the notifications interval. The default value is 60 seconds.  Note that sending notifications is enabled by default.	<b>set antispoof notifications interval</b> <i>interval</i>
12.	Optionally, enable duplicate IP address detection.	<b>set antispoof duplicateIP enable</b>

Table 41-1 lists the commands used to disable or reset anti-spoofing features and to manage the binding table entries.

**Table 41-1 Managing Anti-Spoofing Features**

Task	Command(s)
Disable anti-spoofing globally	<code>set antispoof disable</code> <code>clear antispoof</code>
or Reset all anti-spoofing configuration to default values.	<code>clear antispoof all</code>
Disable sending anti-spoofing notifications.	<code>set antispoof notifications disable</code>
Reset the notification interval to the default of 60 seconds.	<code>clear antispoof notifications interval</code>
Disable duplicate IP address detection.	<code>set antispoof duplicateIP disable</code> <code>clear antispoof duplicateIP</code>
Delete an anti-spoofing port class or clear specific configuration values to their defaults.	<code>clear antispoof class class-index [name] [timeout] [threshold-index thresh-index]</code>
Disable DHCP snooping on the specified port or ports.	<code>clear antispoof dhcp-snooping port-string</code>
Disable DHCP snooping MAC verification on the specified port or ports.	<code>clear antispoof dhcp-snooping mac-verification port-string</code>
Reset the DHCP snooping port mode to untrusted for the specified port or ports.	<code>clear antispoof dhcp-snooping port-mode port-string</code>
Disable dynamic ARP inspection on the specified port or ports.	<code>set antispoof arp-inspection disable port-string</code> <code>clear antispoof arp-inspection port-string</code>
Disable IP source guard on the specified port or ports.	<code>set antispoof ip-inspection disable port-string</code> <code>clear antispoof ip-inspection port-string</code>
Remove an anti-spoofing port class assignment from the specified port or ports.	<code>clear antispoof port-class port-string</code>
Delete an anti-spoofing user source MAC address to source IP address binding from the binding table.	<code>clear antispoof binding {port port-string   mac mac-addr   ip ip-addr}</code>
Reset the anti-spoofing threshold counters to 0 by port, MAC address, or IP address.	<code>clear antispoof counters {port port-string   mac mac-addr   ip ip-addr}</code>

[Table 41-2](#) lists the commands used to display anti-spoofing information.

**Table 41-2 Displaying Anti-Spoofing Information**

Task	Command
Display global anti-spoofing values	<code>show antispoof</code>
Display anti-spoofing class information	<code>show antispoof class [class-index]</code>
Display anti-spoofing port configuration	<code>show antispoof port [port-string] [-interesting]</code>
Display anti-spoofing source MAC address to source IP address bindings	<code>show antispoof binding [port port-string] [mac mac-addr] [ip ip-addr] [all] [-verbose]</code>

**Table 41-2 Displaying Anti-Spoofing Information (continued)**

Task	Command
Display anti-spoofing statistics	<b>show antispoof counters</b> [port <i>port-string</i> ] [mac <i>mac-addr</i> ] [ip <i>ip-addr</i> ] [all] [-verbose]

## Code Example

The following example configures anti-spoofing features on a switch at the edge of the network, with two ports connected to a DHCP server and the rest of the ports connected to users. DHCP snooping is configured on the ports connected to the DHCP server so the binding table will be populated by DHCP snooping.

Two sets of user ports are configured for ARP inspection or IP source guard inspection, but are enabled for inspection only, since the binding table entries are added by DHCP snooping on the DHCP server trusted ports. Also, DHCP snooping MAC verification is enabled on the untrusted user ports.

As part of the configuration:

- Two port classes and timeout, threshold, and action values for those classes are configured.
- DHCP-snooping is enabled on the ports connected to the DHCP server, and they are configured as trusted ports.
- A simple policy profile is created that will drop DHCP server traffic and it is applied to untrusted ports.
- DAI and IP source guard are configured for inspection only on user ports 10 through 40.
- MAC verification is enabled on all user ports. (DHCP snooping must also be enabled on these ports for MAC verification to work.)
- The appropriate port class is assigned to the user ports.
- The notifications interval is changed to 30 seconds.
- Anti-spoofing is enabled globally and duplicate IP address detection is enabled.

This example assumes that quarantine policy profile 3 has previously been configured. Refer to the “Authentication Configuration” chapter in this book for more information about using quarantine profiles and the quarantine agent.

```
7100 Chassis(su)->set antispoof class 1 name DHCP
7100 Chassis(su)->set antispoof class 1 timeout 7200
7100 Chassis(su)->set antispoof class 1 threshold-index 1 threshold-value 1
action syslog trap

7100 Chassis(su)->set antispoof class 2 name "IPSG and DAI"
7100 Chassis(su)->set antispoof class 2 timeout 3600
7100 Chassis(su)->set antispoof class 2 threshold-index 1 threshold-value 1 action
syslog
7100 Chassis(su)->set antispoof class 2 threshold-index 2 threshold-value 2 action
trap
7100 Chassis(su)->set antispoof class 2 threshold-index 3 threshold-value 3
quarantine-profile 3 action quarantine

7100 Chassis(su)->set policy profile 1 name DHCP
```



```
7100 Chassis(su)->set policy rule 1 udpsourceportIP 67 mask 16 drop
7100 Chassis(su)->set policy port ge.2.10-40 1

7100 Chassis(su)->set antispoof dhcp-snooping enable ge.2.2,4
7100 Chassis(su)->set antispoof dhcp-snooping port-mode trusted ge.2.2,4

7100 Chassis(su)->set antispoof arp-inspection inspection-only ge.2.10-40
7100 Chassis(su)->set antispoof ip-inspection inspection-only ge.2.10-40
7100 Chassis(su)->set antispoof dhcp-snooping enable ge.2.10-40
7100 Chassis(su)->set antispoof dhcp-snooping mac-verification enable ge.2.10-40

7100 Chassis(su)->set antispoof port-class 1 ge.2.2,4
7100 Chassis(su)->set antispoof port-class 2 ge.2.10-40

7100 Chassis(su)->set antispoof enable
7100 Chassis(su)->set antispoof notifications interval 30
7100 Chassis(su)->set antispoof duplicateIP enable
```

## RADIUS-Snooping Configuration

This document provides the following information about configuring RADIUS-Snooping on the Extreme Networks 7100-Series platforms.

For information about...	Refer to page...
<a href="#">Using RADIUS-Snooping in Your Network</a>	<a href="#">42-1</a>
<a href="#">Implementing RADIUS-Snooping</a>	<a href="#">42-2</a>
<a href="#">RADIUS-Snooping Overview</a>	<a href="#">42-2</a>
<a href="#">Configuring RADIUS-Snooping</a>	<a href="#">42-5</a>
<a href="#">RADIUS-Snooping Configuration Example</a>	<a href="#">42-8</a>
<a href="#">Terms and Definitions</a>	<a href="#">42-10</a>

### Using RADIUS-Snooping in Your Network

RADIUS-Snooping (RS) is one of the Extreme Networks MultiAuth suite of authentication methods. See [Chapter 43, Authentication Configuration](#) for a detailed discussion of the other authentication methods supported by the 7100-Series platform. RS resides on the distribution-tier switch, allowing for management of any directly connected edge switch that uses the RADIUS protocol to authenticate a network end-station, but does not support the full complement of the Extreme Networks Secure Networks™ capabilities.

The RADIUS client edge-switch initiates an authentication request, by sending a RADIUS request to the RADIUS server that resides upstream of the distribution-tier switch. By investigating the RADIUS request frames, RS can determine the MAC address of the end-user device being authenticated. The network administrator creates a user account on the RADIUS server for the end-user that includes any policy, dynamic VLAN assignment, and other RADIUS and RS attributes for this end-station. By investigating the RADIUS response from the RADIUS server, RS can build a MutiAuth session as though the end-user were directly connected to the distribution-tier device.

Sessions detected by RS function identically to local authenticated sessions from the perspective of the Extreme Networks MultiAuth framework, with the exception that RS can not force a reauthentication event; it can only timeout the session.

RADIUS-Snooping allows the Extreme Networks 7100-Series distribution-tier switch to identify RADIUS exchanges between devices connected to edge switches and apply policy to those devices even when the edge switch is from another vendor and does not support policy. RADIUS-Snooping provides, but is not limited to, the following functionalities:

- RFC 3580 Dynamic VLAN assignment
- Authentication modes support

- Idle and session timeouts support
- Multi-user authentication on a port
- Multi-authentication method support

With RS-enabled on the distribution-tier switch, these Secure Networks capabilities can be configured by the network administrator on an end-user basis.

RADIUS-Snooping accounting is supported.

## Implementing RADIUS-Snooping

RS requires that unencrypted RADIUS request frames, from the edge switch, transit the distribution-tier switch, before proceeding to the up-stream RADIUS server for validation.



**Note:** A router cannot reside between the RADIUS client and the distribution-tier switch enabled for RS. The presence of a router would modify the calling-station ID of the RADIUS request frame that RS depends upon to learn the MAC address of the end-station for this session.

To configure RS on a distribution-tier switch:

- Set the global MultiAuth mode to **multi**
- Set the MultiAuth port mode to **auth-opt** for all ports that are part of the RS configuration
- Globally enable RS on the distribution-tier switch
- Enable RS on all ports over which RADIUS request and response frames will transit
- Optionally change the period RS will wait for a RADIUS response frame from the server
- Populate the RADIUS-Snooping flow table with RS client and RADIUS server combinations
- Optionally enable RADIUS-Snooping accounting

## RADIUS-Snooping Overview

This section provides an overview of RADIUS-Snooping configuration and management.

### RADIUS-Snooping Configuration

#### MultiAuth Configuration

MultiAuth must be enabled if the RADIUS-Snooping configuration involves the authentication of more than a single user on a port. There are two aspects to multiauthentication in a RADIUS-Snooping configuration:

- The global MultiAuth mode must be changed from the default mode of **strict** to **multi**, in order to authenticate multiple downstream users.
- The MultiAuth port mode must be set to **auth-opt** for both upstream (to the RADIUS server) and downstream (to the authenticating switch) ports. Setting global MultiAuth to **multi** sets the default port value from **auth-opt** to **force-auth**. Reset the mode for the affected ports to **auth-opt**.

See the “[MultiAuth Authentication](#)” on page 43-7 for a complete discussion on MultiAuth configuration.

## Enabling RADIUS-Snooping

RS is enabled globally on the distribution-tier switch. It is also enabled on the distribution-tier switch ports directly attached to the edge switch that the RADIUS request frames transit, from the edge switch to the RADIUS server, as well as the ports the response frames transit, from the RADIUS server back to the edge switch.

### Configuring Enabled Port Settings

The number of seconds the firmware waits for a RADIUS response after it successfully snoops a RADIUS request can be set per-port. If you do not set this timeout at the port level, the system level setting is used.

The **authallocated** value specifies the maximum number of RS users per port. You can configure this number of allowed RS users on a per port basis. You should set this **authallocated** value equal to or less than the configured value for the **set multiauth port numusers** command. This value is the maximum number of users per port for all authentication clients. Typically, **authallocated** and **multiauth port numusers** are set to the same value.

## Populating the RADIUS-Snooping Flow Table

The RADIUS-Snooping flow table is a filter that determines which RADIUS server and client combinations will be snooped. If the secret is configured, the response frames are checked for valid MD5 checksum, in order to validate the sender.

The RS flow table contains RADIUS server and client entries for each RADIUS server and client combination for which RS will be used on this system. The RADIUS client IP address and authenticating RADIUS server IP address are manually entered into the RADIUS-Snooping flow table. By default, the RADIUS-Snooping flow table is empty. Entries are added to the flow table based upon an index entry. The first matching entry in the table is used for the continuation of the authentication process.

When an investigated RADIUS frame transits the RS-enabled port with a match in the flow table, RS will track that RADIUS request and response exchange and will build a MultiAuth session for the end-user, based upon what it finds in the RADIUS response frames.

## Setting the RADIUS-Snooping Timeout

A timeout is configured to set the number of seconds that the firmware waits for a RADIUS response frame to be returned from the RADIUS server, after successfully snooping a RADIUS request frame from the client. If no response is seen before the timeout expires, the session is terminated.

## RADIUS-Snooping Management

RADIUS-Snooping management options are available to:

- Terminate all RS sessions or on a per port or MAC address basis
- Reset all RS configuration to its default settings
- Clear all RADIUS-Snooping flow table entries or per index entry
- Display RS statistics

## RADIUS Session Attributes

The RADIUS attributes defining the session are returned in the RADIUS response frame. RADIUS attributes are used to configure the user on the system. Attributes explicitly supported by RS that may be included in the RADIUS response frame are:

- Idle-Timeout – If no frames are seen from this MAC address, for the number of seconds configured, the session will be terminated.
- Session-Timeout – The session is terminated after the number of seconds configured.
- Filter-ID - Defines the policy profile (role) and CLI management privilege level, just as it would for any other local authentication agent.
- Tunnel-Group-Id – Specifies the VLAN ID for this session.



**Note:** Numerous attributes may be supported by the RADIUS client for general RADIUS protocol support. Such attributes are beyond the scope of this document. This RS implementation does not interfere with normal RADIUS client attribute support. The list above indicates attributes actually used by this RADIUS-Snooping application once authentication is successfully completed.

Figure 42-1 RADIUS-Snooping Overview

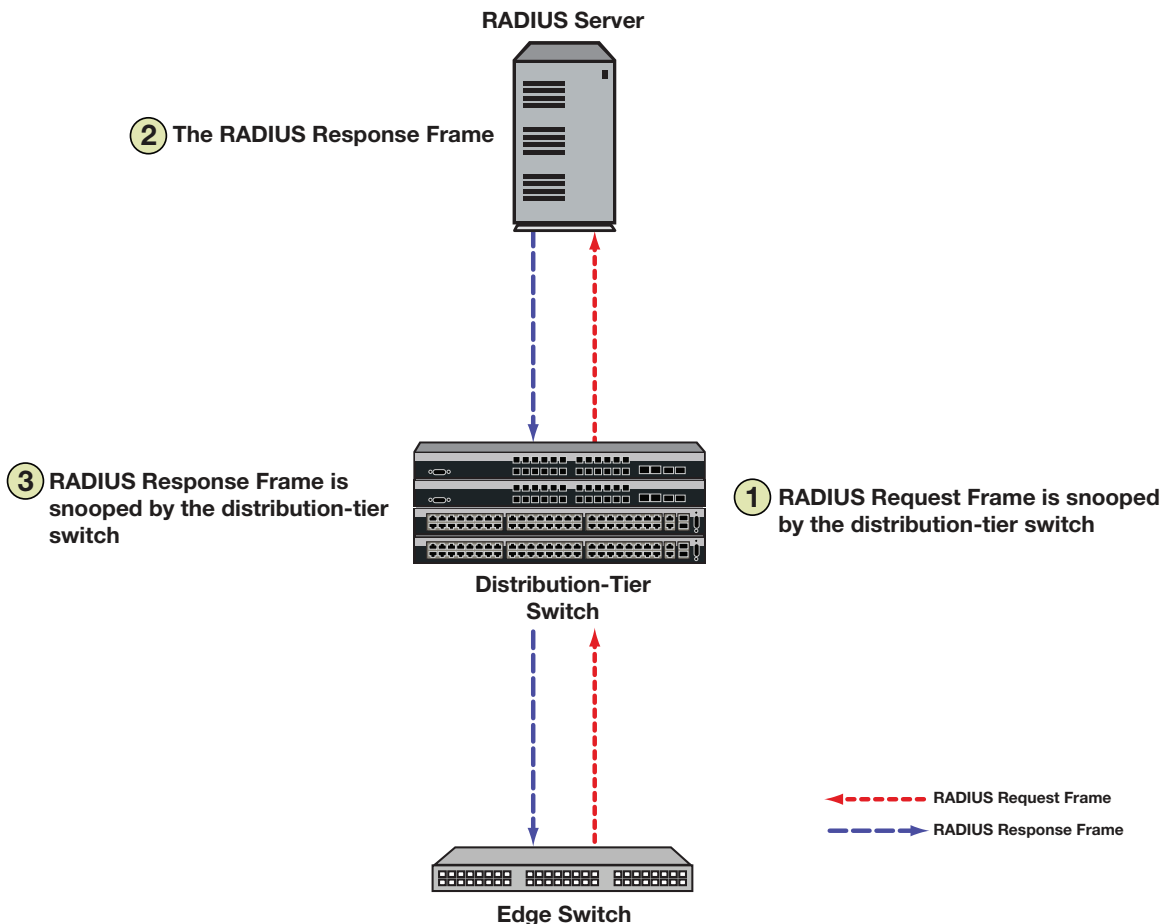


Figure 42-1 illustrates the RADIUS request frame and RADIUS response frame paths. As the RADIUS request frame from the RADIUS client edge device transits the distribution-tier switch, it is snooped. An RS session is created on the distribution-tier switch, if:

- RADIUS snooping is enabled on the switch

- RADIUS-Snooping is enabled on the port
- The RADIUS client edge device and RADIUS server combination are defined in the RADIUS snooping flow table

When the RADIUS server receives the request, the authenticating device is first validated. After validating the authenticating device, the server authenticates the user session itself based on passed username and password attributes. If that succeeds an access accept message containing RADIUS attributes is sent back to the client, otherwise an access reject message is sent back. As the RADIUS response frame transits the distribution-tier switch, the RADIUS attributes contained in the response frame are applied to this session, if an RS session was created for this client server combination and the session has not timed out.

RADIUS-Snooping agent accounting is supported and defaults to disabled. To use RADIUS-Snooping accounting, RADIUS accounting must be enabled using the **set radius accounting** command. RADIUS-Snooping agent accounting can be enabled using the **set radius-snooping accounting** command.

## Configuring RADIUS-Snooping

This section provides details for the configuration of RADIUS-Snooping on the 7100-Series products.

[Table 42-1](#) lists RS parameters and their default values.

**Table 42-1 Default Authentication Parameters**

Parameter	Description	Default Value
RADIUS-Snooping timeout	Specifies the number of seconds that the firmware waits, from the time it successfully snoops a RADIUS request frame, for a RADIUS response frame from the RADIUS server, before terminating the session.	20 seconds
RADIUS-Snooping accounting	Specifies whether RADIUS-Snooping accounting is enabled or disabled on the device.	Disabled
RS system and port state	Enables or disables RS on the distribution-tier switch in a system context or on this port in a port context.	Disabled
authallocated	Specifies the maximum number of allowed RS sessions from all RADIUS clients, on a per port basis.	8
drop	Specifies traffic drop behavior for this port.	Disabled
index	The numeric ID of a RADIUS-Snooping flow table entry.	None
UDP port	Specifies the RADIUS UDP port.	1812
secret	Specifies the RADIUS secret for this RADIUS-Snooping flow table entry.	No secret

## Configuring RADIUS-Snooping on the Distribution-Tier Switch

Procedure 42-1 describes how to configure RADIUS-Snooping on the distribution-tier switch.

### Procedure 42-1 RADIUS-Snooping Configuration

Step	Task	Command(s)
1.	Globally enable MultiAuth for <b>multi</b> mode.	<b>set multiauth mode multi</b>
2.	Configure each upstream and downstream port for the <b>auth-opt</b> mode.	<b>set multiauth port mode auth-opt port-string</b>
3.	Globally enable RADIUS-Snooping on the distribution-tier switch.	<b>set radius-snooping enable</b>
4.	Optionally enable RADIUS-Snooping accounting on the device.	<b>set radius-snooping accounting enable</b>
5.	Enable RADIUS-Snooping on each distribution-tier switch port over which RADIUS request and response frames transit.	<b>set radius-snooping port [enable] [timeout seconds] [authallocated number] [port-string]</b>
6.	Configure RADIUS-Snooping flow table index entries.	<b>set radius-snooping flow index {client-IP-Address server-IP-Address {port} [secret]}</b>
7.	Optionally modify the RADIUS-Snooping timeout setting.	<b>set radius-snooping timeout seconds</b>

## Managing RADIUS-Snooping

Table 42-2 describes how to manage RADIUS-Snooping on the distribution-tier switch.

### Table 42-2 Managing RADIUS-Snooping

Task	Command(s)
To terminate active sessions on the system for the specified port or MAC address.	<b>set radius-snooping initialize {port port-string   mac-address}</b>
To reset all RS configuration to its default value on this system.	<b>clear radius-snooping all</b>
To clear all entries or the specified index entry from the RS flow table.	<b>clear radius-snooping flow {all   index}</b>

## Displaying RADIUS-Snooping Statistics

Table 42-3 describes how to display RADIUS-Snooping statistics.

### Table 42-3 Displaying RADIUS-Snooping Statistics

Task	Command(s)
To display a general overview of the global RS status.	<b>show radius-snooping</b>
To display the RS status for the specified port.	<b>show radius-snooping port port-string</b>
To display information for all or the specified flow index entry.	<b>show radius-snooping flow {index   all}</b>

**Table 42-3 Displaying RADIUS-Snooping Statistics (continued)**

---

To display a summary of sessions for the specified port or MAC address.	<b>show radius-snooping session</b> { <b>port</b> <i>port-string</i>   <b>mac</b> <i>mac-address</i> }
---	--

---

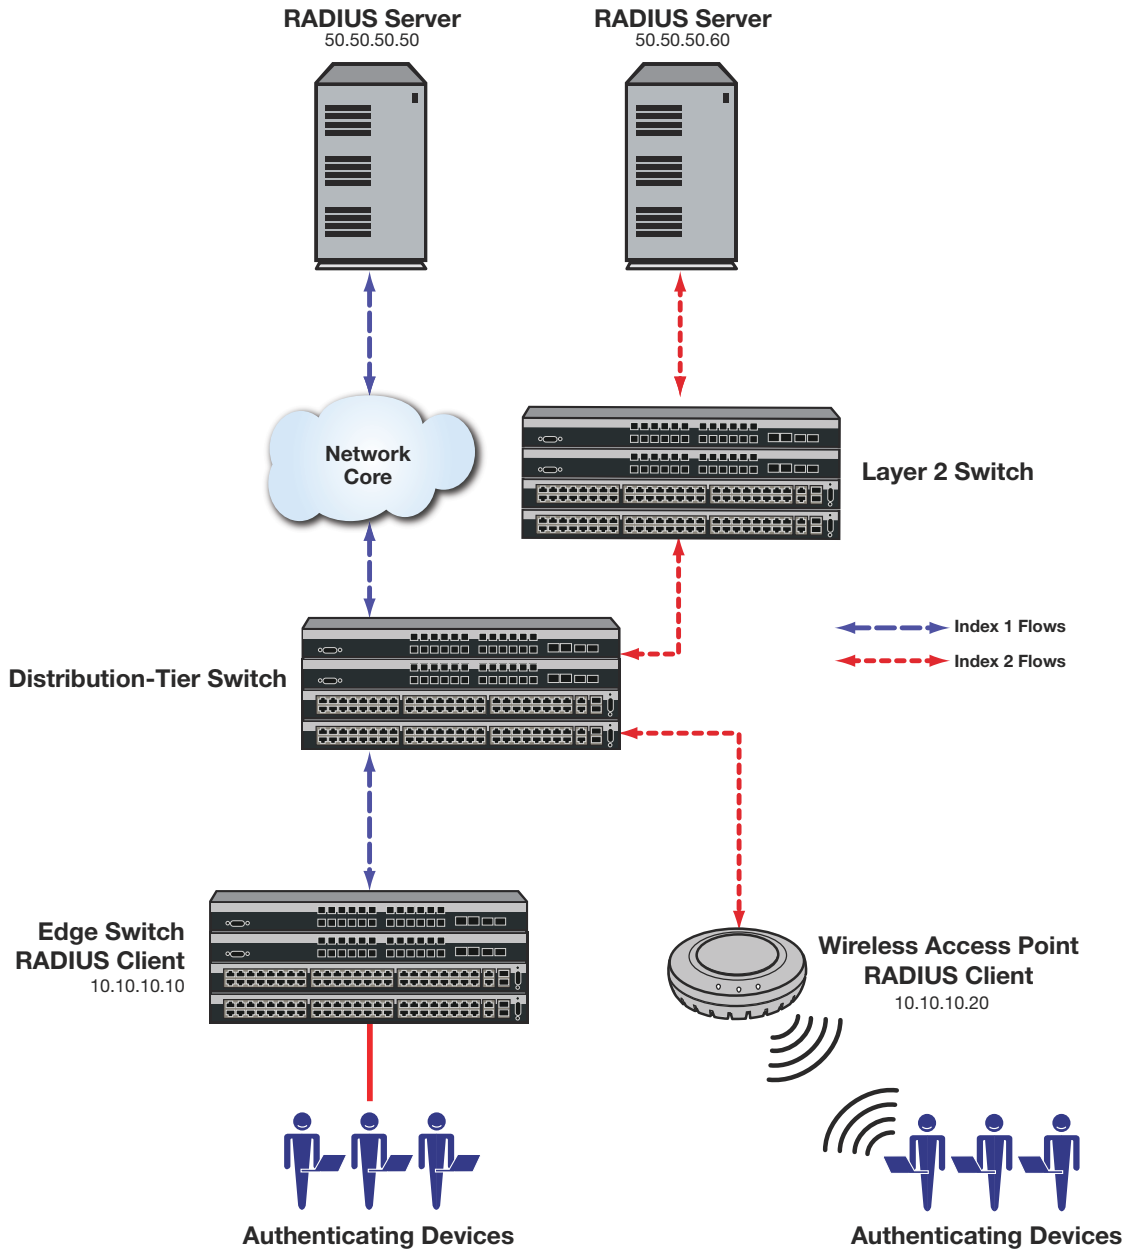


## RADIUS-Snooping Configuration Example

Our RADIUS-Snooping configuration example will configure a distribution-tier switch for two RADIUS request and response flows (index 1 and index 2). Index 1 is from RADIUS client 10.10.10.10 through the network core to the RADIUS server 50.50.50.50. Index 2 is from RADIUS client 10.10.10.20 through a layer 2 switch to the local RADIUS server 50.50.50.60. Each flow is transiting the single distribution-tier switch configured in this example.

See [Figure 42-2](#) for an illustration of the example setup.

**Figure 42-2 RADIUS-Snooping Configuration Example Overview**



We first enable RADIUS-Snooping at the system level for the distribution-tier switch. We then enable two sets of ports (ge.1.5-10 and ge.1.15-24) over which all RADIUS-Snooping request and response frames will transit. In the same command line we:

- Set the maximum number of RS sessions per port to 256

We then configure the two flows as specified above for UDP port 1812 and a secret of **mysecret**.

We complete the configuration by changing the timeout value at the system level to **15** seconds from a default of **20** seconds.

## Configure the Distribution-tier Switch

### Set the MultiAuth mode for the system

```
7100 Chassis(su)->set multiauth mode multi
```

### Set the MultiAuth authentication mode for each port

```
7100 Chassis(su)->set multiauth port mode auth-opt ge.1.5-10,15-24
```

### Enable RS on this system:

```
7100 Chassis(su)->set radius-snooping enable
```

### Enable RS and set configuration for ports on this system

```
7100 Chassis(su)->set radius-snooping port enable authallocated 256 ge.1.5-10
```

```
7100 Chassis(su)->set radius-snooping port enable authallocated 256 ge.1.15-24
```

### Configure RS flow table entries

```
7100 Chassis(su)->set radius-snooping flow 1 10.10.10.10 50.50.50.50 1812 mysecret
```

```
7100 Chassis(su)->set radius-snooping flow 2 10.10.10.20 50.50.50.60 1812 mysecret
```

### Configure RS timeout for this system

```
7100 Chassis(su)->set radius-snooping timeout 15
```

## Managing RADIUS-Snooping on the Distribution-tier Switch

### Terminate an active session on port ge.1.15:

```
7100 Chassis(su)->set radius-snooping initialize port ge.1.15
```

### Reset all RS configuration to its default value:

```
7100 Chassis(su)->clear radius-snooping all
```

### Clear entry index 2 from the RS flow table:

```
7100 Chassis(su)->clear radius-snooping flow 2
```

This completes the RADIUS-Snooping configuration example.

## Terms and Definitions

Table 42-4 lists terms and definitions used in this RADIUS-Snooping configuration discussion.

**Table 42-4 RADIUS-Snooping Configuration Terms and Definitions**

Term	Definition
Calling-Station ID	An attribute field in the RADIUS request and response frames containing the RADIUS client MAC address.
Distribution-Tier Switch	The switch that aggregates edge switch traffic heading into the core network or other distribution devices.
Edge Switch	The switch directly connected to the end-user device.
Filter-ID	A vendor defined RADIUS attribute that the Extreme Networks 7100-Series authentication implementation makes use of, allowing the authenticating device to assign policy, CLI privilege level, and dynamic VLAN assignment to the end-user.
Multi-Authentication Methods	The ability to authenticate a user for multiple authentication methods such as 802.1x, MAC, PWA, or CEP, while only applying the authentication method with the highest authentication precedence.
Multi-User Authentication	The ability to authenticate multiple users on a port, assigning unique policy to each user based upon the user account RADIUS server configuration and policy configuration on the distribution-tier switch.
MutiAuth Framework	The aspect of Secure Networks functionality that provides authentication capabilities including, but not limited to, multi-user and multi-method authentication, application of policy and Dynamic VLAN assignment.
RADIUS Client	In a RADIUS-Snooping context the RADIUS client is the non-Secure Networks capable edge switch that is responsible for authenticating its attached end-user device or port.
RADIUS-Snooping flow table	A table containing the RADIUS client and server ID defining valid RS sessions.
RADIUS Request Frames	Frames sent by the RADIUS client to the RADIUS server requesting end-user authentication validation.
RADIUS Response Frames	Frames sent by the RADIUS server to the RADIUS client either validating or rejecting an authentication validation request. These frames can also contain the Filter-ID attribute allowing the assignment of policy, CLI privilege, and dynamic VLAN assignment.
RADIUS-Snooping	Provides non-Secure Networks capable edge switches with the full range of Secure Networks authentication capabilities when the RADIUS server is upstream of the distribution-tier switch.

## Authentication Configuration

This document provides the following information about configuring user authentication on the Extreme Networks 7100-Series platforms.

For information about...	Refer to page...
<a href="#">Using Authentication in Your Network</a>	43-1
<a href="#">Implementing User Authentication</a>	43-2
<a href="#">Authentication Overview</a>	43-2
<a href="#">Configuring Authentication</a>	43-15
<a href="#">Authentication Configuration Example</a>	43-31
<a href="#">Terms and Definitions</a>	43-36

### Using Authentication in Your Network

Authentication is the ability of a network access server, with a database of valid users and devices, to acquire and verify the appropriate credentials of a user or device (supplicant) attempting to gain access to the network. Extreme Networks authentication uses the RADIUS protocol to control access to switch ports from an authentication server and to manage the message exchange between the authenticating device and the server. Both MultiAuth and Multi-User authentication are supported. MultiAuth is the ability to configure multiple authentication modes for a user and apply the authentication mode with the highest precedence. Multi-User is the ability to appropriately authenticate multiple supplicants on a single link and provision network resources, based upon an appropriate policy for each supplicant. The Extreme Networks switch products support the following seven authentication methods:

- Quarantine agent
- IEEE 802.1x
- Port Web Authentication (PWA)
- MAC-based Authentication (MAC)
- Convergence End Point (CEP)
- RADIUS Snooping (Refer to [Chapter 42, RADIUS-Snooping Configuration](#) for RADIUS Snooping configuration details)
- Auto tracking

Extreme Networks switch products support the configuration of up to seven simultaneous authentication methods per user, with a single authentication method applied based upon MultiAuth authentication precedence.

Network resources represent a major capital investment for your organization and can be vulnerable to both undesired resource usage and malicious intent from outside users. Authentication provides you with a user validation function which assures that the supplicant requesting access has the right to do so and is a known entity. To the degree a supplicant is not a known entity, access can be denied or granted on a limited basis. The ability of authentication to both validate a user's identity and define the resources available to the user assures that valuable network resources are being used for the purposes intended by the network administrator.

Accounting is supported for all authentication agent types.

## Implementing User Authentication

Take the following steps to implement user authentication:

- Determine the types of devices to be authenticated.
- Determine the correct authentication type for each device.
- Determine an appropriate policy best suited for the use of that device on your network.
- Configure RADIUS user accounts on the authentication server for each device.
- Configure user authentication.

## Authentication Overview

For information about...	Refer to page...
<a href="#">Quarantine</a>	43-2
<a href="#">IEEE 802.1x Using EAP</a>	43-3
<a href="#">MAC-Based Authentication (MAC)</a>	43-4
<a href="#">Port Web Authentication (PWA)</a>	43-4
<a href="#">Convergence End Point (CEP)</a>	43-5
<a href="#">Auto-Tracking</a>	43-5
<a href="#">Multi-User And MultiAuth Authentication</a>	43-6
<a href="#">Remote Authentication Dial-In Service (RADIUS)</a>	43-9

### Quarantine

The quarantine agent works in conjunction with a quarantine policy rule to perform the action specified in the associated policy role if the policy rule is hit. The quarantine agent also acts in conjunction with anti-spoofing and will perform the configured class action if an anti-spoofing class threshold is met (see [Chapter 41, Anti-Spoofing Configuration](#) for anti-spoofing configuration details).

The quarantine agent must be enabled globally on the switch and locally on the port to be operational on the port. The quarantine agent is a form of authentication that depends upon the existence of one or more configured quarantine policy rules, with each rule associated with a policy profile. To configure a policy rule as a quarantine profile, configure the policy rule with the desired traffic filtering specifications and specify the quarantine-profile rule option, indicating the associated policy profile.

Once one or more quarantine policy rules are configured and associated with a policy profile, the quarantine authentication agent behaves as any other MultiAuth authentication agent. By default, the quarantine agent has the highest configurable MultiAuth precedence. Static rules have the highest multiauth precedence. Static rule multiauth precedence is not configurable.

There are two circumstance for which actions specified in a quarantine policy are used:

- A quarantine policy rule is hit. In this case, the quarantine agent becomes one of the authentication agents from which the authentication provisioning result will be chosen based upon MultiAuth precedence. So long as the default precedence is not changed, if a quarantine policy rule hit occurs, quarantine authentication is selected and any actions configured in the policy profile taken.
- An anti-spoofing class threshold has been met for which a quarantine action has been configured.

Should you configure quarantine authentication for a lower MultiAuth precedence, if a non-quarantine authentication agent both returns a result and has the highest MultiAuth precedence, quarantine authentication will not be used in that context. If you change the quarantine agent MultiAuth precedence level to a lower precedence, make sure this is the behavior you want.

Quarantine agent accounting is supported and defaults to disabled. To use quarantine agent accounting, RADIUS accounting must be enabled using the **set radius accounting** command. Quarantine agent accounting can be enabled using the **set quarantine-agent accounting** command.

## IEEE 802.1x Using EAP

The IEEE 802.1x port-based access control standard allows you to authenticate and authorize user access to the network at the port level. Access to the switch ports is centrally controlled from an authentication server using RADIUS. The Extensible Authentication Protocol (EAP), defined in RFC 3748, provides the means for communicating the authentication information.

There are three supported types of EAP:

- **MD5** – EAP-MD5 is a challenge-handshake protocol over EAP that authenticates the user with a normal username and password.
- **TLS** – EAP-TLS provides a transport layer security based upon the presentation and acceptance of digital certificates between the supplicant and the authentication server.
- **Protected** – Protected Extensible Authentication Protocol (PEAP) optionally authenticates the authentication server to the client using an X-509 certificate using a TLS tunnel, after which the client authentication credentials are exchanged.

All Extreme Networks platforms support IEEE 802.1x, which protects against unauthorized access to a network, DoS attacks, theft of services and defacement of corporate web pages.

802.1x configuration consists of setting port, global 802.1x parameters, and RADIUS parameters on the switches to point the switch to the authentication server. The Filter-ID RADIUS attribute can be configured on the authentication server to direct dynamic policy assignment on the switch to the 802.1x authenticating end system.

802.1x agent accounting is supported and defaults to enabled. RADIUS accounting must be enabled using the **set radius accounting** command. 802.1x agent accounting can be enabled using the **set dot1x accounting** command.

## MAC-Based Authentication (MAC)

MAC-based authentication (MAC) authenticates a device using the source MAC address of received packets. Two modes are supported for MAC authentication: password and RADIUS username. The MAC authentication mode is set using the **set macauthentication auth-mode** command.

By default the MAC authentication server uses an administratively configured password to authenticate a user. The default value for the password is "NOPASSWORD". The administratively configured password is set using the **set macauthentication password** command.

MAC authentication can be configured to use the RADIUS server configured username credential where the password is the same as the username. The following is an example RADIUS server configuration for MAC address 00-00-22-22-02-01, first with a mask of 48, followed by the address with a mask of 40.

### The full user name with a mask of 48:

```
00-00-22-22-02-01 Auth-Type := Local, User-Password == "00-00-22-22-02-01"  
Service-Type = Framed-User
```

### The user name with a mask of 40:

```
00-00-22-22-03-00 Auth-Type := Local, User-Password == "00-00-22-22-03-00"  
Service-Type = Framed-User
```

In either case, if the authentication server receives valid credentials from the switch, RADIUS returns an Accept message to the switch.

MAC authentication enables switches to authenticate end systems, such as printers and camcorder devices that do not support 802.1x or web authentication. Since MAC-based authentication authenticates the device, not the user, and is subject to MAC address spoofing attacks, it should not be considered a secure authentication method. However, it does provide a level of authentication for a device where otherwise none would be possible.

MAC-based authentication agent accounting is supported and defaults to enabled. RADIUS accounting must be enabled using the **set radius accounting** command. MAC-based authentication agent accounting can be disabled using the **set macauthentication accounting** command.

## Port Web Authentication (PWA)

Port Web Authentication (PWA) authenticates a user by utilizing a web browser for the login process to authenticate to the network. To log in using PWA, a user opens the web browser requesting a URL that either directly accesses the PWA login page or is automatically redirected to the login page. At the PWA login page, the user enters a login username and password. On the switch, either the Challenge Handshake Authentication Protocol (CHAP) or the Password Authentication Protocol (PAP) verifies the username and password credentials provided to the authentication server. If the credentials are validated, the authentication server returns a RADIUS Access-Accept message, optionally containing Filter-ID or tunnel attributes, to the switch.

PAP uses an unencrypted password. CHAP uses the password to generate a digest that is transmitted to the authentication server. If RADIUS determines that the digest matches the digest generated on the authentication server, access is granted. The acceptance message back to the switch can contain any Filter-ID attribute configured on the authentication server, allowing policy to be applied for the authenticating user.

PWA agent accounting is supported and defaults to enabled. RADIUS accounting must be enabled using the **set radius accounting** command. PWA agent accounting can be disabled using the **set pwa accounting** command.

## Convergence End Point (CEP)

CEP detects an IP telephony or video device on a port and dynamically applies a specific policy to the port. The switch detects a convergence end point by inspecting received packets for specific traffic attributes. CEP does not require a RADIUS configuration.

The CEP implementation supports the following detection methods:

- **Cisco Phone Detection** - the firmware parses a Cisco Discovery Protocol (CDP) packet to identify the phone type. If it was sent by an IP phone, the firmware uses the phone type. A response is sent back to the phone, verifying authentication.
- **Siemens HiPath Phone Detection** - TCP/UDP port number snooping is used. Port 4060 is the default port for communication.
- **H.323 Phone Detection** - TCP/UDP port number snooping and reserved IP address snooping are used. Ports 1718 - 1720 and IP address 224.0.1.41 are the default values.
- **LLDP-MED** - Detects LLDP-MED on the specified port.
- **Session Initiation Protocol (SIP) Phone Detection** - TCP/UDP port number snooping and reserved IP address snooping are used. Port 5060 and IP address 224.0.1.75 are the default values.

CEP agent accounting is supported and defaults to disabled. To use CEP agent accounting RADIUS accounting must be enabled using the **set radius accounting** command. PWA agent accounting can be enabled using the **set pwa accounting** command.

## Auto-Tracking

The auto-tracking agent is a form of authentication that authenticates those sessions that are not captured by the other supported MultiAuth authentication agents (quarantine, 802.1x, PWA, MAC, CEP, and RADIUS snooping). If auto-tracking is disabled, these sessions are never entered into the session table. Many policy driven switch features depend on the session being in the session table for the feature to interact with the session. It is important that a network administrator have the ability to determine which station addresses on which ports are not being authenticated through traditional MultiAuth methods. Auto-tracking provides the administrator with the ability to assign these sessions a provisioning result based upon the contents of the admin-policy. Because these sessions can now be tracked, an administrator can determine whether and how to provision them in the future, allowing for increased security and control.

The auto-tracking authentication agent must be enabled globally on the switch and locally on the port to be operational on the port.

The auto-tracking authentication agent behaves the same as any other authentication agent, with the exception that it always returns an authentication result. By default, the auto-tracking agent has the lowest MultiAuth precedence. The auto-tracking agent is one of the authentication agents from which the authentication provisioning result will be chosen based upon MultiAuth precedence. Each authentication agent attempts to authenticate the user. All authentication agents that return a result are grouped. The authentication agent with the highest MultiAuth precedence is selected to authorize the user. For the default MultiAuth precedence ordering, all other authentication agents must fail to return an authentication result for auto-tracking to be selected. If auto-tracking is the selected authentication method, an auto-tracking session is created and if an admin-policy exists, the admin-policy provisions the user session.

It is recommended that you do not configure auto-tracking authentication for a higher MultiAuth precedence than its default setting of lowest. If a non-auto-tracking authentication agent both returns a result and has a lower MultiAuth precedence, that authentication method will never be used, because auto-tracking always returns a result and has been configured with a higher MultiAuth precedence.



Auto-tracking agent accounting is supported and defaults to disabled. To use auto-tracking accounting, RADIUS accounting must be enabled using the **set radius accounting** command. Auto-tracking agent accounting can be enabled using the **set pwa accounting** command.

Auto-tracking can be configured with a RADIUS timeout profile. The RADIUS timeout profile allows you to provision a session that encounters a RADIUS timeout condition, on a per port basis, with a policy profile other than the default policy. The RADIUS timeout profile allows a MAC address that attempted to authenticate during a RADIUS outage to be dealt with in a non-default manner based upon the contents of the specified policy profile. The RADIUS timeout profile is configured using the **set auto-tracking port radius-timeout-profile** command.

Auto-tracking can be configured with a RADIUS access reject profile. The RADIUS access reject profile allows you to provision a session that encounters a RADIUS access reject response from the RADIUS server, on a per port basis, with a policy profile other than the default policy. The RADIUS access reject profile allows a MAC address that was rejected by the RADIUS server to be dealt with in a non-default manner based upon the contents of the specified policy profile.

The RADIUS access reject profile takes precedence over the RADIUS timeout profile configured using the **set auto-tracking port radius-timeout-profile** command, should a RADIUS timeout take place and a RADIUS access reject has already occurred for this session.

The RADIUS access reject profile is configured using the **set auto-tracking port radius-reject-profile** command.

## Multi-User And MultiAuth Authentication

This section will discuss multi-user and MultiAuth authentication. Multi-user and MultiAuth are separate concepts. The primary difference between the two is as follows:

- Multi-user authentication refers to the ability to authenticate multiple users and devices on the same port, with each user or device being provided the appropriate level of network resources based upon policy.
- MultiAuth authentication refers to the ability of a single or multiple user(s), device(s), or port(s) to successfully authenticate using multiple authentication methods at the same time, such as 802.1x, PWA, and MAC, with precedence determining which authentication method is actually applied to that user, device, or port.

### Multi-User Authentication

Multi-user authentication provides for the per-user or per-device provisioning of network resources when authenticating. It supports the ability to receive from the authentication server:

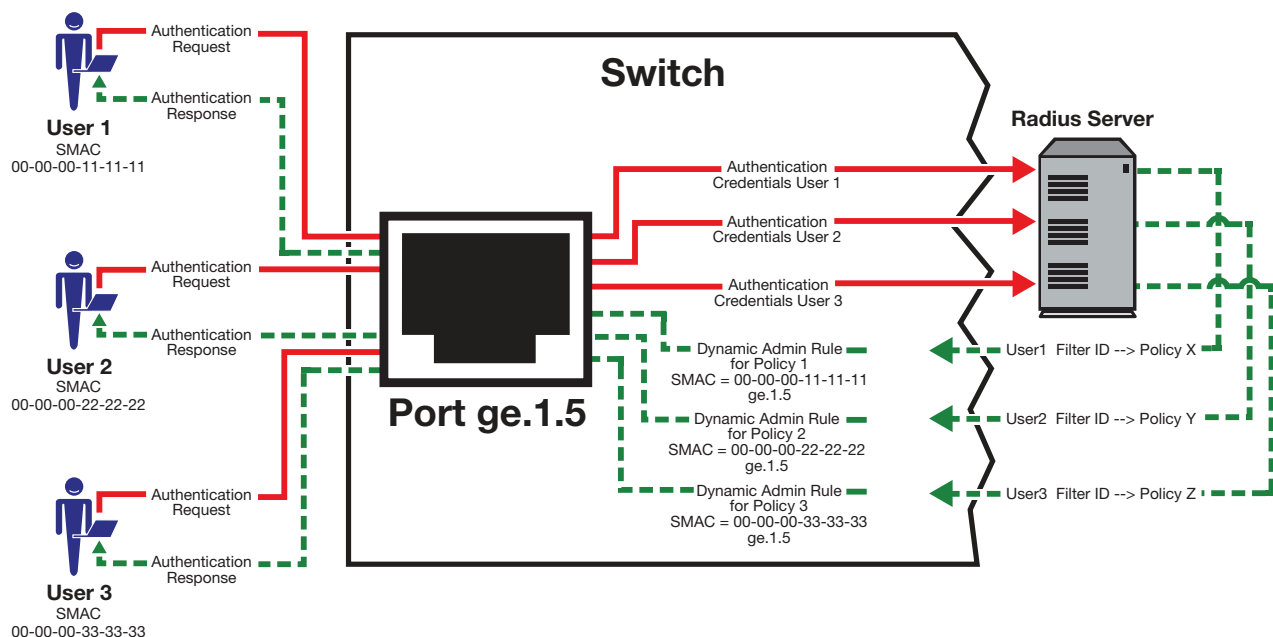
- A policy traffic profile, based on the user account's RADIUS Filter-ID configuration
- A base VLAN-ID, based on the RFC 3580 tunnel attributes configuration, also known as dynamic VLAN assignment

When a single supplicant connected to an access layer port authenticates, a policy profile can be dynamically applied to all traffic on the port. When multi-user authentication is not implemented, and more than one supplicant is connected to a port, firmware does not provision network resources on a per-user or per-device basis. Different users or devices may require a different set of network resources. The firmware tracks the source MAC address for each authenticating user regardless of the authenticating protocol being used. Provisioning network resources on a per-user basis is accomplished by applying the policy configured in the RADIUS Filter-ID, or the base VLAN-ID configured in the RFC 3580 tunnel attributes, for a given user's MAC address. The RADIUS Filter-ID and tunnel attributes are part of the RADIUS user account and are included in the RADIUS Access-Accept message response from the authentication server.

The number of allowed users per port can be configured using the **set multiauth port numusers** command. See the **set multiauth port** command in the *Extreme Networks 7100-Series CLI Reference* for the number of supported users per module. The **show multiauth port** command displays both the allowed number of users configured and the maximum number of users supported per port for the device. The allowed number of users defaults to the maximum number of supported users for the port.

In [Figure 43-1](#) each user on port ge.1.5 sends an authentication request to the RADIUS server. Based upon the Source MAC address (SMAC), RADIUS looks up the account for that user and includes the Filter-ID associated with that account in the authentication response back to the switch (see section “[The RADIUS Filter-ID](#)” on page 43-10 for Filter-ID information). The policy specified in the Filter-ID is then applied to the user. See section “[RFC 3580](#)” on page 43-11 for information on dynamic VLAN assignment and tunnel attribute configuration.

**Figure 43-1 Applying Policy to Multiple Users on a Single Port**



## MultiAuth Authentication

Authentication mode support provides for the global setting of a single authentication mode 802.1X (strict-mode) or multiple modes (MultiAuth) per user or port when authenticating.

Strict mode is the appropriate mode when authenticating a single 802.1X user. All traffic on the port receives the same policy in strict mode. When authenticating PWA, CEP, or MAC, you must use MultiAuth authentication, whether authenticating a single or multiple supplicants.

MultiAuth authentication supports the simultaneous configuration of up to seven authentication methods per user on the same port, but only one method per user is actually applied. When MultiAuth authentication ports have a combination of authentication methods enabled, and a user is successfully authenticated for more than one method at the same time, the configured authentication method precedence will determine:

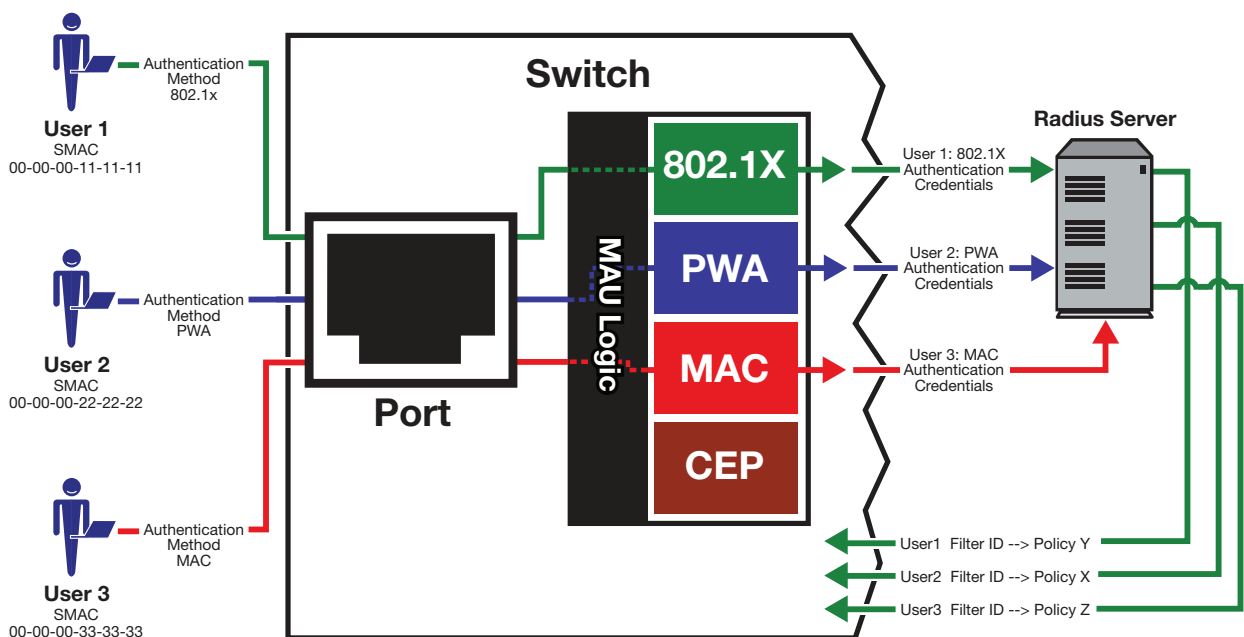
- Which RADIUS-returned filter ID will be processed and result in an applied traffic policy profile, in the case of 802.1X, MAC, PWA, and CEP. See “[Setting MultiAuth Authentication Precedence](#)” on page 43-24 for authentication method precedence details.
- Whether a quarantine policy is applied, in the case of the quarantine agent

- Whether an auto-tracking session is created and an admin-policy (if it exists) is applied

The number of users or devices MultiAuth authentication supports depends upon the type of device, whether the ports are fixed access or uplink, and whether increased port capacity or extra chassis user capacity MUA licenses have been applied. See the firmware customer release note that comes with your device for details on the number of users or devices supported per port.

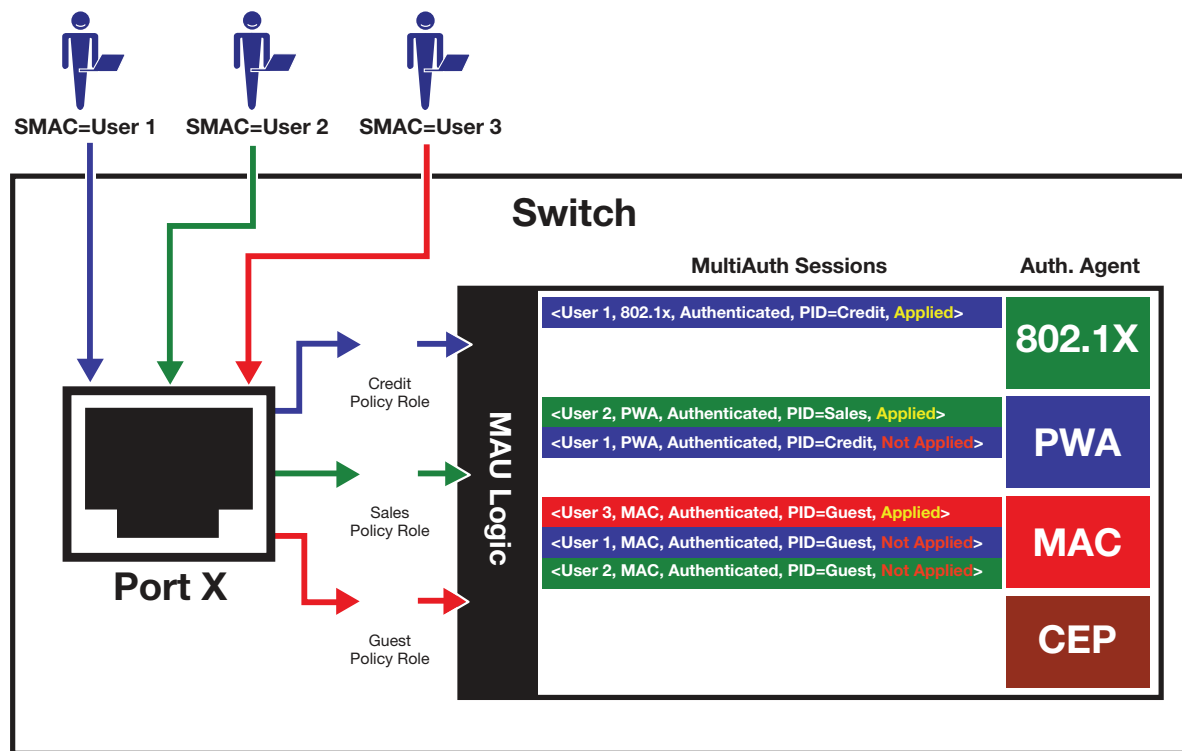
In Figure 43-2, multiple users are authenticated on a single port each with a different authentication method (in this example only 802.1X, PWA, MAC, and CEP are enabled on the device). In this case, each user on a single port successfully authenticates with a different authentication type. The authentication method is included in the authentication credentials sent to the RADIUS server. RADIUS looks up the user account for that user based upon the SMAC. The filter ID for that user is returned to the switch in the authentication response, and the authentication is validated for that user.

**Figure 43-2 Authenticating Multiple Users With Different Methods on a Single Port**



In Figure 43-3, full MultiAuth authentication takes place in that multiple users on a single port are validated for more than one authentication method. The applied authentication and policy are based upon the authentication method precedence level. On the far right column of the figure, the enabled authentication methods are listed from top to bottom in order of precedence. User 1 is authenticating with both the 802.1x and PWA methods, with the Credit policy. Both the 802.1x and PWA authentication methods are validated, but only the 802.1x MultiAuth session is applied, because that has the highest precedence. User 2 is authenticating with both PWA and MAC methods, with the Sales policy. PWA, having a higher precedence than MAC, is the MultiAuth session applied for User 2. User 3 is a guest and is authenticating with the MAC method only. The MAC MultiAuth session, with the Guest policy is applied for User 3.

Figure 43-3 Selecting Authentication Method When Multiple Methods are Validated



When a re-authentication attempt times out, the timeout action can either be set to terminate the session or for none, in which case the session remains authenticated and provisioned according to the prior successful RADIUS authentication response. It is recommended that you not set the re-authentication timeout action to none when using 802.1x authentication.

## Remote Authentication Dial-In Service (RADIUS)

This section provides details for the configuration of RADIUS and RFC 3580 attributes.

For information about...	Refer to page...
<a href="#">How RADIUS Data Is Used</a>	<a href="#">43-10</a>
<a href="#">The RADIUS Filter-ID</a>	<a href="#">43-10</a>
<a href="#">RADIUS Authentication Retransmission Algorithm</a>	<a href="#">43-11</a>
<a href="#">RFC 3580</a>	<a href="#">43-11</a>
<a href="#">Policy Mappable Response</a>	<a href="#">43-13</a>

The Remote Authentication Dial-In User Service (RADIUS) is an extensible protocol used to carry authentication and authorization information between the switch and the Authentication Server (AS). RADIUS is used by the switch for communicating supplicant supplied credentials to the authentication server and the authentication response from the authentication server back to the switch. This information exchange occurs over the link-layer protocol.

The switch acts as a client to RADIUS using UDP port 1812 by default (configurable in the **set radius** command). The authentication server contains a database of valid supplicant user accounts with their corresponding credentials. The authentication server checks that the information received from the switch is correct, using authentication schemes such as PAP, CHAP, or EAP. The

authentication server returns an Accept or Reject message to the switch based on the credential validation performed by RADIUS. The implementation provides enhanced network security by using a shared secret and MD5 password encryption.

Required authentication credentials depend upon the authentication method being used. For 802.1x and PWA authentication, the switch sends username and password credentials to the authentication server. For MAC authentication, the switch sends the device MAC address and a password configured on the switch to the authentication server. The authentication server verifies the credentials and returns an Accept or Reject message back to the switch.

## How RADIUS Data Is Used

The Extreme Networks switch bases its decision to open the port and apply a policy or close the port based on the RADIUS message, the port's default policy, and unauthenticated behavior configuration.

RADIUS provides accounting functionality by way of accounting packets from the switch to the RADIUS server, for such session statistics as start and end, total packets, and session end reason events. This data can be used for both billing and network monitoring purposes.

Additionally RADIUS is widely used by VoIP service providers. It is used to pass login credentials of a SIP end point (like a broadband phone) to a SIP Registrar using digest authentication, and then to the authentication server using RADIUS. Sometimes it is also used to collect call detail records (CDRs) later used, for instance, to bill customers for international long distance.

If you configure an authentication method that requires communication with an authentication server, you can use the RADIUS Filter-ID attribute to dynamically assign either a policy profile or management level to authenticating supplicants.

## The RADIUS Filter-ID

The RADIUS Filter-ID attribute consists of a string that is formatted in the RADIUS Access-Accept packet sent back from the authentication server to the switch during the authentication process.

Each user can be configured in the RADIUS server database with a RADIUS Filter-ID attribute that specifies the name of either a policy profile or management level the user should be assigned upon successful authentication. During the authentication process, when the authentication server returns a RADIUS Access-Accept packet that includes a Filter-ID matching a policy profile name configured on the switch, the switch then dynamically applies the policy profile to the physical port the supplicant is authenticating on.

The decorated Filter-ID supports a policy attribute, a management access attribute, or both in the following formats:

```
Enterasys:version=1:policy=polycyname
```

```
Enterasys:version=1:mgmt=access-mgmtType
```

```
Enterasys:version=1:mgmt=access-mgmtType:policy=polycyname
```

*polycyname* is the name of the policy to apply to this authentication.

*access-mgmtTypes* supported are: **ro** (read-only), **rw** (read-write), and **su** (super-user).

The un-decorated Filter-ID supports the policy attribute only in the following format:

```
polycyname
```

The undecorated format is simply a string that specifies a policy profile name. The undecorated format cannot be used for management access authentication. Decorated Filter-IDs are processed first. If no decorated Filter-IDs are found, then undecorated Filter-IDs are processed. If multiple Filter-IDs are found that contain conflicting values, a Syslog message is generated.

## RADIUS Authentication Retransmission Algorithm

There are three RADIUS authentication algorithms:

**Standard** – RADIUS authentication always uses the primary (lowest server ID) RADIUS server if it is reachable. If a network outage occurs or server capacity is exceeded, secondary RADIUS servers are used. The standard RADIUS authentication algorithm is appropriate when multiple RADIUS servers are used for redundancy as opposed to a scaled provisioning environment.

**Round Robin** – RADIUS authentications are evenly spread across servers, allowing the load balancing of a large number of authentications across all available RADIUS servers. If a given server goes down, only sessions associated with that server are affected.

**Sticky Round Robin** – RADIUS attempts to use the same RADIUS server for any given authentication session, but uses round robin assigning a RADIUS server to each unique authentication session. The sticky round robin algorithm is appropriate for devices that support a limited number of sessions such as the Extreme Network Access Controller (NAC).

The RADIUS authentication algorithm setting defaults to standard. Use the **set radius algorithm** command to set RADIUS authentication algorithm globally on the device.

## RADIUS Authentication Sticky Round Robin Maximum Sessions

Round robin sessions are only associated with a particular server when using the sticky round robin algorithm. The maximum number of sticky round robin sessions allowed can be configured on a per port basis. The maximum sessions setting can be configured between 0 (no sessions allowed) and the maximum number of users allowed on the system as displayed using the **show multiauth** command. The RADIUS authentication algorithm must be set to sticky round robin using the **set radius algorithm** command.

If **max-sessions** is not specified, the maximum number of sticky round robin sessions supported defaults to the maximum number of users supported on the device as displayed in the show multiauth command.

Use the **set radius max-sessions** command, specifying the **max-sessions** value and the server index, to set the maximum allowed sessions on when the RADIUS authentication algorithm is set to sticky round robin.

## RFC 3580

Extreme Networks switches support the RFC 3580 RADIUS tunnel attribute for dynamic VLAN assignment. The VLAN-Tunnel-Attribute implements the provisioning of service in response to a successful authentication. On ports that do not support policy, the packet will be tagged with the VLAN-ID. The VLAN-Tunnel-Attribute defines the base VLAN-ID to be applied to the user.

### Dynamic VLAN Assignment

The RADIUS server may optionally include RADIUS tunnel attributes in a RADIUS Access-Accept message for dynamic VLAN assignment of the authenticated end system.

RFC 3580's RADIUS tunnel attributes are often configured on a RADIUS server to dynamically assign users belonging to the same organizational group within an enterprise to the same VLAN, or to place all offending users according to the organization's security policy in a Quarantine VLAN. Tunnel attributes are deployed for enterprises that have end system authentication configured on the network. For example, all engineers can be dynamically assigned to the same VLAN upon authentication, while sales are assigned to another VLAN upon authentication.

The name of the feature on Extreme Networks platforms that implements dynamic VLAN assignment through the receipt of RADIUS tunnel attributes is VLAN authorization. VLAN authorization depends upon receipt of the RFC 3580 RADIUS tunnel attributes in RADIUS Access-Accept messages. VLAN authorization must be enabled globally and on a per-port basis

for the Tunnel attributes to be processed. When disabled per port or globally, the device will not process Tunnel attributes.

By default, all policy-capable Extreme Networks platforms will dynamically assign a policy profile to the port of an authenticating user based on the receipt of the Filter-ID RADIUS attribute. This is not the case for RADIUS tunnel attributes in that, by default, VLAN authorization is disabled.

### **VLAN Authorization Attributes**

Three Tunnel attributes are used for dynamic VLAN Authorization:

- Tunnel-Type attribute (Type=64, Length=6, Tag=0, Value=0x0D for VLAN)
- Tunnel-Medium-Type attribute (Type=65, Length=6, Tag=0, Value=0x06 for 802 media)
- Tunnel-Private-Group-ID attribute (Type=81, Length>=3, String=VID in ASCII)

The Tunnel-Type attribute indicates the tunneling protocol to be used when this attribute is formatted in RADIUS Access-Request messages, or the tunnel protocol in use when this attribute is formatted in RADIUS Access-Accept messages. Set Tunnel-Type attribute parameters as follows:

- Type: Set to 64 for Tunnel-Type RADIUS attribute
- Length: Set to 6 for six-byte length of this RADIUS attribute
- Tag: Provides a means of grouping attributes in the same packet which refer to the same tunnel. Valid values for this field are from 0x01 through 0x1F, inclusive. Set to 0 if unused. Unless alternative tunnel types are provided, it is only necessary for tunnel attributes to specify a single tunnel. As a result, where it is only desired to specify the VLAN-ID, the tag field should be set to zero (0x00) in all tunnel attributes.
- Value: Indicates the type of tunnel. A value of 0x0D (decimal 13) indicates that the tunneling protocol is a VLAN.

Tunnel-Medium-Type indicates the transport medium to use when creating a tunnel for the tunneling protocol, determined from Tunnel-Type attribute. Set Tunnel-Medium-Type attribute parameters as follows:

- Type: Set to 65 for Tunnel-Medium-Type RADIUS attribute
- Length: Set to 6 for six-byte length of this RADIUS attribute
- Tag: Provides a means of grouping attributes in the same packet which refer to the same tunnel. Valid value for this field are 0x01 through 0x1F, inclusive. Set to 0 if unused. Unless alternative tunnel types are provided, it is only necessary for tunnel attributes to specify a single tunnel. As a result, where it is only desired to specify the VLANID, the tag field should be set to zero (0x00) in all tunnel attributes.
- Value: Indicates the type of tunnel. A value of 0x06 indicates that the tunneling medium pertains to 802 media (including Ethernet)

Tunnel-Private-Group-ID attribute indicates the group ID for a particular tunneled session. Set the Tunnel-Private-Group-ID attribute parameters as follows:

- Type: Set to 81 for Tunnel-Private-Group-ID RADIUS attribute
- Length: Set to a value greater than or equal to 3.
- Tag: Provides a means of grouping attributes in the same packet which refer to the same tunnel. Valid values for this field are from 0x01 through 0x1F, inclusive. Set to 0 if unused. Unless alternative tunnel types are provided, it is only necessary for tunnel attributes to specify a single tunnel. As a result, where it is only desired to specify the VLANID, the tag field should be set to zero (0x00) in all tunnel attributes.

- **String:** Indicates the group. For the VLAN ID integer value, it is encoded as a string using ASCII. For example, the VLAN ID integer value 103 would be represented as 0x313033

## VLAN Authorization Considerations

VLAN Authorization poses some operational and management issues on the network.

- A VLAN is not a security container. It is a broadcast container and used to segment broadcast traffic on the network. ACLs implemented at the layer 3 routed interface for a VLAN only provide access control for traffic into and out of the VLAN. No access control mechanism for intra-VLAN communications exists, therefore users within the VLAN are not protected from each other. Malicious traffic allowed onto a VLAN can potentially infect all traffic on the VLAN. Such an infection can consume valuable hardware resources on the infrastructure, such as CPU cycles and memory. Infections can be transmitted to other hosts within the VLAN and to the layer 3 routed boundary. This leads to the direct competition of malicious traffic with business critical traffic on the network.
- End-To-End QoS cannot be truly guaranteed if QoS is implemented at the layer 3 routed interface for a network where business critical applications are classified and prioritized.
- If VLANs are implemented to group together users that are members of the same organizational group, then a VLAN must be configured everywhere in the network topology where a member of that organizational unit may connect to the network. For example, if an engineer may connect to the network from any location, then the Engineering VLAN must be configured on all access layer devices in the network. These VLAN configurations lead to over-extended broadcast domains as well as added configuration complexity in the network topology.
- A problem with moving an end system to a new VLAN is that the end system must be issued an IP address on the new VLAN's subnet to which it has become a member. If the end system does not yet have an IP address, this is not usually a problem. However, if the end system has an IP address, the lease of the address must time out before it attempts to obtain a new address, which may take some time. The IP address assignment process, implemented by DHCP, and the authentication process are not conjoined on the end system. Therefore, this leads to end systems possessing an invalid IP address after dynamic VLAN Authorization and lost IP connectivity until its current IP address times out. Furthermore, when a new IP address is eventually assigned to the end system, IP connectivity is disrupted for all applications on the end system.

## Policy Mappable Response

The policy mappable response, or conflict resolution, feature allows you to define how the system should handle allowing an authenticated user onto a port based on the contents of the RADIUS Accept message reply. There are three possible response settings: tunnel mode, policy mode, or both tunnel and policy, also known as hybrid authentication mode.

When the mappable response is set to **tunnel** mode, the system will use the tunnel attributes in the RADIUS reply to apply a VLAN to the authenticating user and will ignore any Filter-ID attributes in the RADIUS reply. When tunnel mode is configured, VLAN-to-policy mapping can occur.

When the mappable response is set to **policy** mode, the system will use the Filter-ID attributes in the RADIUS reply to apply a policy to the authenticating user and will ignore any tunnel attributes in the RADIUS reply. When policy mode is configured, no VLAN-to-policy mapping will occur.

When the mappable response is set to **both**, or hybrid authentication mode, both Filter-ID attributes (dynamic policy assignment) and tunnel attributes (dynamic VLAN assignment) sent in RADIUS Accept message replies are used to determine how the switch should handle authenticating users. When hybrid authentication mode is configured, VLAN-to-policy mapping can occur, as described below in [When Policy Mappable Response is "Both"](#).



Using hybrid authentication mode eliminates the dependency on having to assign VLANs through policy roles — VLANs can be assigned by means of the tunnel attributes while policy roles can be assigned by means of the Filter-ID attributes. Alternatively, VLAN-to-policy mapping can be used to map policies to users using the VLAN specified by the tunnel attributes, without having to configure Filter-ID attributes on the RADIUS server. This separation gives administrators more flexibility in segmenting their networks beyond the platform's policy role limits.

### When Policy Mappable Response is “Both”

Hybrid authentication mode uses both Filter-ID attributes and tunnel attributes. To enable hybrid authentication mode, use the **set policy mappable** command and set the **response** parameter to **both**. When configured to use both sets of attributes:

- If both the Filter-ID and tunnel attributes are present in the RADIUS reply, then the policy profile specified by the Filter-ID is applied to the authenticating user, and if VLAN authorization is enabled globally and on the authenticating user's port, the VLAN specified by the tunnel attributes is applied to the authenticating user.

If VLAN authorization is not enabled, the VLAN specified by the policy profile is applied. See [“RFC 3580”](#) on page 43-11 for information about VLAN authorization.

- If the Filter-ID attributes are present but the tunnel attributes are not present, the policy profile specified by the Filter-ID is applied, along with the VLAN specified by the policy profile.
- If the tunnel attributes are present but the Filter-ID attributes are not present, and if VLAN authorization is enabled globally and on the authenticating user's port, then the switch will check the VLAN-to-policy mapping table (configured with the **set policy mappable** command):
  - If an entry mapping the received VLAN ID to a policy profile is found, then that policy profile, along with the VLAN specified by the policy profile, will be applied to the authenticating user.
  - If no matching mapping table entry is found, the VLAN specified by the tunnel attributes will be applied to the authenticating user.
  - If the VLAN-to-policy mapping table is invalid, then the `etsysPolicyRFC3580MapInvalidMapping` MIB is incremented and the VLAN specified by the tunnel attributes will be applied to the authenticating user.

If VLAN authorization is not enabled, the tunnel attributes are ignored.

### When Policy Mappable Response is “Profile”

When the switch is configured to use only Filter-ID attributes, by setting the **set policy mappable** command **response** parameter to **policy**:

- If the Filter-ID attributes are present, the specified policy profile will be applied to the authenticating user. If no Filter-ID attributes are present, the default policy (if it exists) will be applied.
- If the tunnel attributes are present, they are ignored. No VLAN-to-policy mapping will occur.

### When Policy Mappable Response is “Tunnel”

When the switch is configured to use only tunnel attributes, by setting the **set policy mappable** command **response** parameter to **tunnel**, and if VLAN authorization is enabled both globally and on the authenticating user's port:

- If the tunnel attributes are present, the specified VLAN will be applied to the authenticating user. VLAN-to-policy mapping can occur.

- If the tunnel attributes are not present, the default policy VLAN will be applied; if the default policy VLAN is not configured, the port VLAN will be applied.
- If the Filter-ID attributes are present, they are ignored.

If VLAN authorization is not enabled, the user will be allowed onto the port with the default policy, if it exists. If no default policy exists, the port VLAN will be applied.

## Configuring Authentication

This section provides details for the configuration of authentication methods, MultiAuth and RADIUS.

For information about...	Refer to page...
<a href="#">Configuring Quarantine Agent</a>	<a href="#">43-17</a>
<a href="#">Configuring IEEE 802.1x</a>	<a href="#">43-18</a>
<a href="#">Configuring MAC-based Authentication</a>	<a href="#">43-19</a>
<a href="#">Configuring Port Web Authentication (PWA)</a>	<a href="#">43-20</a>
<a href="#">Configuring Convergence End Point (CEP)</a>	<a href="#">43-21</a>
<a href="#">Configuring Auto-Tracking</a>	<a href="#">43-22</a>
<a href="#">Configuring MultiAuth Authentication</a>	<a href="#">43-23</a>
<a href="#">Configuring RADIUS</a>	<a href="#">43-28</a>

Table 43-1 lists Authentication parameters and their default values.

**Table 43-1 Default Authentication Parameters**

Parameter	Description	Default Value
auto-tracking agent authallocated	Specifies the maximum number of users per port supported by the auto-tracking agent	Number of MultiAuth users configured.
auto-tracking agent port idle timeout	Specifies the auto-tracking agent port idle timeout value in seconds	MultiAuth port idle timeout.
auto-tracking agent port session timeout	Specifies the auto-tracking agent port session timeout in seconds.	MultiAuth port session timeout.
auto-tracking agent state	Enables or disables the auto-tracking agent on a per port basis and globally on the device.	Disabled.
authentication agent accounting	Specifies whether accounting is enabled or disabled for the agent.	Quarantine – Disabled. 802.1x – Enabled. PWA – Enabled. MAC-based authentication – Enabled. CEP – Disabled. Auto-tracking – Disabled.
cep port	Enables or disables CEP for the specified port.	Disabled.

**Table 43-1 Default Authentication Parameters (continued)**

Parameter	Description	Default Value
dot1x	Enables and disables 802.1x authentication both globally and per port.	Globally: Disabled. Per Port: Enabled.
dot1x authconfig	Configures 802.1x authentication.	auto - auto authorization mode.
macauthentication	Globally enables or disables MAC authentication on a device.	Disabled.
macauthentication authallocated	Sets the number of MAC authentication sessions supported on the specified port.	8.
macauthentication port	Enables or disables MAC authentication on a port	Disabled.
MultiAuth idle-timeout	Specifies the period length for which no traffic is received before a MultiAuth session is set to idle.	300 seconds.
MultiAuth mode	Globally sets MultiAuth for this device.	strict - authentication limited to 802.1x for a single user on a port.
MultiAuth port mode	Specifies the MultiAuth port mode to use for the specified port.	auth-opt - Authentication is optional based upon global and port configuration.
MultiAuth precedence	Specifies the authentication mode to use when multiple authentication types are successfully authenticated.	Precedence from high to low: 802.1x, PWA, MAC, CEP, Radius-Snooping.
MultiAuth session-timeout	Specifies the maximum amount of time a session can live.	0 - no timeout in effect.
pwa	Globally enables or disables PWA authentication.	Disabled.
quarantine agent authallocated	Specifies the maximum number of users per port for supported by the quarantine agent	Number of Multiauth users configured.
quarantine agent port idle timeout	Specifies the quarantine agent port idle timeout value in seconds	Multiauth port idle timeout.
quarantine agent port session timeout	Specifies the quarantine agent port session timeout in seconds.	Multiauth port session timeout.
quarantine agent state	Enables or disables the quarantine agent on a per port basis and globally on the device.	Disabled.
radius	Enable or disable RADIUS on this device.	Disabled.
radius accounting	Enables or disables RADIUS accounting for this device.	Disabled.
radius accounting intervalminimum	Specifies the minimum interval before sending updates for RADIUS accounting.	600 seconds.

**Table 43-1 Default Authentication Parameters (continued)**

Parameter	Description	Default Value
radius accounting retries	Specifies the number of times a switch will attempt to contact an authentication server for RADIUS accounting that is not responding.	2.
radius accounting timeout	Specifies the amount of time for a switch to make contact with a RADIUS server.	5 seconds.
radius accounting updateinterval	Specifies the minimum interval between interim updates for RADIUS accounting.	1800 seconds.
radius authentication algorithm	The algorithm used for selecting the server used for a RADIUS authentication session.	standard.
radius retries	Specifies the number of times a switch will try to establish with the authentication server.	3.
RADIUS sticky round robin maximum sessions	The maximum number of RADIUS authentication sessions allowed when the RADIUS authentication algorithm is set to sticky round robin.	maximum number of users supported on the device.
radius timeout	Specifies the amount of time a switch will wait to receive a response from the authentication server before sending another request.	20 seconds.
realm	Specifies authentication server configuration scope	Both: management-access and network-access.
VLAN authorization	Enables or disables globally and per port VLAN authorization.	Globally: Disabled. Per Port: Enabled.
VLAN egress format	Determines whether dynamic VLAN tagging will be none, tagged, untagged, or dynamic for an egress frame.	Untagged.

## Configuring Quarantine Agent

Configuring the Quarantine agent on an authenticator switch port consists of:

- Setting the quarantine agent state globally and per port
- Setting optional quarantine agent parameters

[Procedure 43-1](#) describes how to configure the quarantine agent on an authenticator switch port. Unspecified parameters use their default values.

### Procedure 43-1 Quarantine Agent Configuration

Step	Task	Command(s)
1.	Enable the quarantine agent on the switch.	<b>set quarantine-agent enable</b>
2.	Enable the quarantine agent on each port.	<b>set quarantine-agent port enable <i>port-string</i></b>

**Procedure 43-1 Quarantine Agent Configuration (continued)**

Step	Task	Command(s)
3.	Optionally, modify the maximum number of quarantine agent sessions allowed on the port. Defaults to number of Multiauth users setting.	<b>set quarantine-agent port authallocated</b> <i>num-users port-string</i>
4.	Optionally, modify the idle timeout value in seconds. Defaults to the Multiauth idle timeout value.	<b>set quarantine-agent port idle-timeout</b> <i>idle-timeout port-string</i>
5.	Optionally, modify the session timeout value in seconds. defaults to the global multiauth session timeout.	<b>set quarantine-agent port session-timeout</b> <i>session-timeout port-string</i>
6.	Optionally, globally enable quarantine agent accounting.	<b>set quarantine-agent accounting {enable   disable}</b>
7.	Optionally display global quarantine agent state and the ports with quarantine agent enabled.	<b>show quarantine-agent [port port-string]</b>

## Configuring IEEE 802.1x

Configuring IEEE 802.1x on an authenticator switch port consists of:

- Setting the authentication mode globally and per port
- Configuring optional authentication port parameters globally and per port
- Globally enabling 802.1x authentication for the switch

[Procedure 43-2](#) describes how to configure IEEE 802.1x on an authenticator switch port. Unspecified parameters use their default values.

**Procedure 43-2 IEEE 802.1x Configuration**

Step	Task	Command(s)
1.	Set the IEEE 802.1x authentication mode both globally and per port: <ul style="list-style-type: none"> <li>• <b>Auto</b> - The switch will only forward authenticated frames.</li> <li>• <b>Forced-auth</b> - 802.1x authentication is effectively disabled for this port. All received frames are forwarded.</li> <li>• <b>Forced-unauth</b> - 802.1x authentication is effectively disabled on the port. If 802.1x is the only authentication method on the port, all frames are dropped.</li> </ul>	<b>set dot1x auth-config</b> <b>authcontrolled-portcontrol {auto   forced-auth   forced-unauth}</b>
<p><b>Note:</b> Before enabling 802.1x authentication on the switch, you must set the authentication mode of ports that will not be participating in 802.1x authentication to forced-authorized to assure that frames will be forwarded on these ports. Examples of this kind of port are connections between switches and connections between a switch and a router.</p> <p>See the <i>Extreme Networks 7100-Series CLI Reference</i> for a listing of parameter options that come with this command.</p>		

**Procedure 43-2 IEEE 802.1x Configuration (continued)**

Step	Task	Command(s)
2.	Display the access entity index values. Ports used to authenticate and authorize supplicants utilize access entities that maintain entity state, counters, and statistics for an individual supplicant. You need to know the index value associated with a single entity to enable, disable, initialize, or reauthenticate a single entity.	<b>show dot1x auth-session-stats</b>
3.	Enable IEEE 802.1x globally on the switch. Ports default to enabled.	<b>set dot1x {enable   disable} [port-string] [index index-list]</b>
4.	If an entity deactivates due to the supplicant logging off, inability to authenticate, or the supplicant or associated policy settings are no longer valid, you can reinitialize a deactivated access entity. If necessary, reinitialize the specified entity.	<b>set dot1x init [index index-list]</b>
5.	If the authentication for a supplicant times out or is lost for any reason, you can reauthenticate that supplicant. If necessary, reauthenticate the specified entity.	<b>set dot1x reauth [index index-list]</b>
6.	Optionally, globally disable 802.1x agent accounting.	<b>set dot1x accounting {enable   disable}</b>
7.	Display IEEE 802.1x configuration.	<b>show dot1x auth-config</b>

## Configuring MAC-based Authentication

Configuring MAC-based authentication on a switch consists of:

- Setting the global MAC authentication password for the switch
- Optionally setting the number of MAC authentication sessions allowed on a port
- Enabling MAC authentication on a port
- Enabling MAC authentication globally
- Setting the authentication mode to multi
- Optionally reinitializing or reauthenticating existing sessions

[Procedure 43-3](#) describes how to configure MAC-based authentication. Unspecified parameters use their default values.

**Procedure 43-3 MAC-Based Authentication Configuration**

Step	Task	Command(s)
1.	Optionally set or clear a global password on the switch.	<b>set macauthentication password password</b> <b>clear macauthentication password password</b>
2.	Set or clear the number of MAC authentication sessions supported on a port.	<b>set macauthentication authallocated number port-string</b>

**Procedure 43-3 MAC-Based Authentication Configuration (continued)**

Step	Task	Command(s)
3.	Enable or disable MAC authentication on a port. By default, MAC authentication is disabled for all ports. MAC authentication must be enabled on the ports that will use it.	<b>set macauthentication port {enable   disable}</b>
4.	Set the authentication mode for the credentials sent to the authentication server.	<b>set macauthentication auth-mode {password   radius-username}</b>
5.	Enable or disable MAC authentication globally on the device. By default, MAC authentication is globally disabled on the device.	<b>set macauthentication {enable   disable}</b>
6.	Set the MultiAuth mode.	<b>set multiauth mode multi</b>
7.	Optionally, globally disable MAC-based authentication accounting.	<b>set macauthentication accounting {enable   disable}</b>
8.	Display MAC authentication configuration or status of active sessions.	<b>show macauthentication</b> <b>show macauthentication session</b>
9.	If a session or port requires reinitialization, reinitialize a specific MAC session or port.	<b>set macauthentication macinitialize</b> <i>mac-address</i> <b>set macauthentication portinitialize</b> <i>port-string</i>
10.	If a session or port requires reauthentication, reauthenticate a specific MAC session or port.	<b>set macauthentication macreauthenticate</b> <i>mac-address</i> <b>set macauthentication portreauthenticate</b> <i>port-string</i>

## Configuring Port Web Authentication (PWA)

Configuring PWA on the switch consists of:

- Setting the IP address which the user will authenticate to on the switch
- Enabling PWA on the port
- Globally enabling PWA on the switch
- Setting the authentication mode

[Procedure 43-4](#) describes how to configure PWA authentication. Unspecified parameters use their default values.

**Procedure 43-4 Port Web Authentication (PWA) Configuration**

Step	Task	Command(s)
1.	Set the IP address for the end-station the supplicant accesses.	<b>set pwa ipaddress</b> <i>ip-address</i>
2.	Enable or disable PWA. PWA must be enabled on the port for PWA to function.	<b>set pwa portcontrol enable</b> <i>port-string</i> <b>set pwa portcontrol disable</b> <i>port-string</i>
3.	Globally enable or disable PWA on the switch.	<b>set pwa enable</b> <b>set pwa disabled</b>
4.	Optionally, globally disable PWA authentication accounting.	<b>set pwa accounting {enable   disable}</b>

**Procedure 43-4 Port Web Authentication (PWA) Configuration (continued)**

Step	Task	Command(s)
5.	Set the MultiAuth mode.	<b>set multiauth mode multi</b>
6.	Display PWA configuration.	<b>show pwa</b>

## Configuring Convergence End Point (CEP)

Configuring CEP consists of:

- Creating a CEP detection group for Non-Cisco Detection CEP types
- Enabling the CEP group for Cisco Detection
- Setting the CEP policy per CEP type
- Enabling CEP on the port
- Setting the authentication mode

### Creating a CEP Detection Group

CEP detection groups can be created, deleted, enabled, or disabled. You create a CEP detection group by associating an ID with the create command. Once a group is created, you associate a CEP type, IP address, protocol, and high or low protocol port to it. The type can be H.323, Siemens, or SIP. The IP address is the IP address of the CEP device. By default, H.323 will use 224.0.1.41 as its IP address and Siemens will have no IP address configured. The protocol can be TCP or UDP. The high or low protocol port is the maximum or minimum TCP or UDP port to be used by the group.

[Procedure 43-5](#) describes the creation of a CEP detection group.

**Procedure 43-5 CEP Detection Group Configuration**

Step	Task	Command(s)
1.	Create a new CEP detection group or enable, disable, or delete an existing group.	<b>set cep detection-id id {create   enable   disable   delete}</b>
2.	Specify the CEP type to be associated with the this group.	<b>set cep detection-id id type {h323   siemens   sip}</b>
3.	Specify the CEP device IP address and mask or set to unknown.	<b>set cep detection-id id address {ip-address   unknown} mask {mask   unknown}</b>
4.	Set the CEP detection group protocol.	<b>set cep detection-id id protocol {tcp   udp   both   none}</b>
5.	Set the maximum or minimum port for the TCP or UDP group protocol.	<b>set cep detection-id id {porthigh   portlow} port</b>

[Procedure 43-6](#) describes the steps to configure CEP.

**Procedure 43-6 CEP Configuration**

Step	Task	Command(s)
1.	Determine the policy profile index of the profile you wish to associate with a CEP type.	<b>show policy profile all</b>
2.	Associate a policy profile with a CEP type.	<b>set cep policy {cisco   h323   lldp-med   siemens   sip} policy-index</b>



**Procedure 43-6 CEP Configuration (continued)**

Step	Task	Command(s)
3.	Enable or disable the CEP device port for the CEP type	<b>set cep port <i>port-string</i> <i>cep-type</i> enable</b> <b>set cep port <i>port-string</i> <i>cep-type</i> disable</b>
4.	If you are using the Cisco discovery protocol, enable the Cisco discovery protocol. You can also optionally set the voice VLAN ID, whether tagged traffic is trusted or untrusted, and 802.1X priority transmitted to the Cisco IP phone to format in the 802.1Q VLAN tag of its VoIP traffic.	<b>set ciscodp port { [status {disable   enable}] [vvid {<i>vlan-id</i>   none   dot1p   untagged}] [trust-ext {trusted   untrusted}] [cos-ext <i>value</i>] } <i>port-string</i></b>
5.	If the Cisco discovery protocol is enabled on any port, enable the Cisco discovery protocol globally.	<b>set ciscodp status</b>
6.	Globally enable or disable CEP on the switch.	<b>set cep enable</b> <b>set cep disable</b>
7.	Optionally, globally enable CEP agent accounting.	<b>set cep accounting {enable   disable}</b>
8.	Set the MultiAuth mode.	<b>set multiauth mode multi</b>
9.	Display CEP connections, detection, policy and port settings.	<b>show cep {connections   detection   policy   port}</b>

**Setting MultiAuth Idle and Session Timeout for CEP**

There is no means of detecting if a Siemens, SIP, or H323 phone goes away other than in the case of a link down. Therefore, if these types of phones are not directly connected to the switch port and the phone goes away, the switch will still see the phone connection and any configured policy will remain on the port. Detected CEPs will be removed from the connection table if they do not send traffic for a time equal to the MultiAuth authentication idle timeout value. CEPs are also removed if the total duration of the session exceeds the time specified in the MultiAuth authentication session timeout.

[Procedure 43-7](#) describes setting the MultiAuth idle and session timeout for CEP.

**Procedure 43-7 Configuring MultiAuth Idle and Session Timeouts for CEP**

Step	Task	Command(s)
1.	Optionally set the MultiAuth authentication idle timeout for this switch.	<b>set multiauth idle-timeout cep <i>timeout</i></b>
2.	Optionally set the MultiAuth authentication session timeout for this switch.	<b>set multiauth session-timeout cep <i>timeout</i></b>

**Configuring Auto-Tracking**

Configuring the auto-tracking agent on an authenticator switch port consists of:

- Setting the auto-tracking agent state globally and per port
- Setting optional auto-tracking agent parameters

[Procedure 43-8](#) describes how to configure the auto-tracking agent on an authenticator switch port. Unspecified parameters use their default values.

**Procedure 43-8 Auto-tracking Agent Configuration**

Step	Task	Command(s)
1.	Enable the auto-tracking agent on the switch.	<b>set auto-tracking enable</b>
2.	Enable the auto-tracking agent on each port.	<b>set auto-tracking port enable</b> <i>port-string</i>
3.	Optionally, modify the maximum number of auto-tracking agent sessions allowed on the port. Defaults to number of Multiauth users setting.	<b>set auto-tracking port authallocated</b> <i>num-users port-string</i>
4.	Optionally, modify the idle timeout value in seconds. Defaults to the Multiauth idle timeout value.	<b>set auto-tracking port idle-timeout</b> <i>idle-timeout port-string</i>
5.	Optionally, modify the session timeout value in seconds. defaults to the global multiauth session timeout.	<b>set auto-tracking port session-timeout</b> <i>session-timeout port-string</i>
6.	Optionally, globally enable auto-tracking agent accounting.	<b>set auto-tracking accounting</b> {enable   disable}
7.	Optionally, configure a RADIUS timeout profile on a per port basis.	<b>set auto-tracking port radius-timeout-profile</b> <i>profile-id port-string</i>
8.	Optionally, configure a RADIUS access reject profile on a per port basis.	<b>set auto-tracking port radius-reject-profile</b> <i>profile-id port-string</i>
9.	Optionally display global auto-tracking agent state and the ports with the auto-tracking agent enabled.	<b>show auto-tracking</b> [ <b>port</b> <i>port-string</i> ]

## Configuring MultiAuth Authentication

For information about...	Refer to page...
<a href="#">Setting MultiAuth Authentication Mode</a>	<a href="#">43-23</a>
<a href="#">Setting MultiAuth Authentication Precedence</a>	<a href="#">43-24</a>
<a href="#">Setting MultiAuth Authentication Port Properties</a>	<a href="#">43-24</a>
<a href="#">Setting MultiAuth Authentication Timers</a>	<a href="#">43-25</a>
<a href="#">Setting MultiAuth Authentication Traps</a>	<a href="#">43-26</a>
<a href="#">Setting the MultiAuth Re-Authentication Timeout Action</a>	<a href="#">43-26</a>
<a href="#">Displaying MultiAuth Configuration Information</a>	<a href="#">43-27</a>
<a href="#">Configuring VLAN Authorization</a>	<a href="#">43-27</a>

### Setting MultiAuth Authentication Mode

MultiAuth authentication mode can be set to MultiAuth or strict 802.1X single user mode. Set MultiAuth authentication to MultiAuth when multiple users need to be authenticated for 802.1X or in all cases for quarantine agent, MAC, PWA, CEP, and auto-tracking agent authentication.

[Procedure 43-9](#) describes setting the MultiAuth authentication mode.

**Procedure 43-9 MultiAuth Authentication Configuration**

Step	Task	Command(s)
1.	For a single user, single authentication 802.1x port configuration, set MultiAuth mode to strict.	<b>set multiauth mode strict</b>
2.	For multiple user 802.1x authentication or any non-802.1x authentication, set the system authentication mode to use multiple authenticators simultaneously.	<b>set multiauth mode multi</b>
3.	To clear the MultiAuth authentication mode.	<b>clear multiauth mode</b>

**Setting MultiAuth Authentication Precedence**

MultiAuth authentication administrative precedence globally determines which authentication method will be selected when a user is successfully authenticated for multiple authentication methods on a single port. When a user successfully authenticates more than one method at the same time, the precedence of the authentication methods will determine which RADIUS-returned filter ID will be processed and result in an applied traffic policy profile.

MultiAuth authentication precedence defaults to the following order from high to low: quarantine agent, 802.1x, PWA, MAC, CEP, Radius-Snooping, auto tracking. You may change the precedence for one or more methods by setting the authentication methods in the order of precedence from high to low. Any methods not entered are given a lower precedence than the methods entered in their pre-existing order. For instance, if you start with the default order and only set quarantine agent, PWA and MAC, the new precedence order will be quarantine agent, PWA, MAC, 802.1x, CEP, and auto-tracking.



**Note:** It is highly recommended that if you are using the quarantine agent authentication method that it always have the highest precedence. It is also highly recommended that you keep the auto tracking authentication method at the lowest precedence.

Given the default order of precedence (quarantine, 802.1x, PWA, MAC, CEP, and auto-tracking), if a user was to successfully authenticate with PWA and MAC, the authentication method RADIUS Filter-ID applied would be PWA, because it has a higher position in the order. A MAC session would authenticate, but its associated RADIUS Filter-ID would not be applied. If no other authentication method successfully authenticated, the auto-tracking agent would authenticate and an auto-tracking session initiated. The session would authenticate based upon the contents of the admin-policy, if an admin-policy exists.

[Procedure 43-10](#) describes setting the order for MultiAuth authentication precedence.

**Procedure 43-10 MultiAuth Authentication Precedence Configuration**

Step	Task	Command(s)
1.	Set a new order of precedence for the selection of the RADIUS filter ID that will be returned when multiple authentication methods are authenticated at the same time for a single user.	<b>set multiauth precedence</b> <b>{[quarantine-agent] [dot1x] [pwa] [mac]</b> <b>[cep] [radius-snooping] [auto-tracking]}</b>
2.	Reset the order MultiAuth authentication precedence to the default values.	<b>clear multiauth precedence</b>

**Setting MultiAuth Authentication Port Properties**

MultiAuth authentication supports the configuration of MultiAuth port and maximum number of users per port properties. The MultiAuth port property can be configured as follows:

- **Authentication Optional** – Authentication methods are active on the port based upon the global and port authentication method. Before authentication succeeds, the current policy role applied to the port is assigned to the ingress traffic. This is the default role if no authenticated user or device exists on the port. After authentication succeeds, the user or device is allowed to access the network according to the policy information returned from the authentication server, in the form of the RADIUS Filter-ID attribute, or the static configuration on the switch. This is the default setting.
- **Authentication Required** – Authentication methods are active on the port, based on the global and per port authentication method configured. Before authentication succeeds, no traffic is forwarded onto the network. After authentication succeeds, the user or device gains access to the network based upon the policy information returned by the authentication server in the form of the RADIUS Filter-ID attribute, or the static configuration on the switch.
- **Force Authenticated** – The port is completely accessible by all users and devices connected to the port, all authentication methods are inactive on the port, and all frames are forwarded onto the network.
- **Force Unauthenticated** – The port is completely closed for access by all users and devices connected to the port. All authentication methods are inactive and all frames are discarded.

[Procedure 43-11](#) describes setting the MultiAuth authentication port and maximum user properties.

#### Procedure 43-11 MultiAuth Authentication Port and Maximum User Properties Configuration

Step	Task	Command(s)
1.	Set the specified ports to the MultiAuth authentication optional port mode.	<b>set multiauth port mode auth-opt</b> <i>port-string</i>
2.	Set the specified ports to the MultiAuth authentication required port mode.	<b>set multiauth port mode auth-reqd</b> <i>port-string</i>
3.	Set the specified ports to the MultiAuth authentication force authenticated port mode.	<b>set multiauth port mode force-auth</b> <i>port-string</i>
4.	Set the specified ports to the MultiAuth authentication force unauthenticated port mode.	<b>set multiauth port mode force-unauth</b> <i>port-string</i>
5.	Optionally set the maximum number of authenticated users for the specified port.  <b>Notes:</b> This value can be set to any value up to the maximum number of MultiAuth users supported for the device. See the firmware release notes that come with your device for the maximum number of supported MultiAuth users the device supports.	<b>set multiauth port mode numusers</b> <i>numusers port-string</i>
6.	Reset the ports MultiAuth authentication port mode to the default value for the specified ports.	<b>clear multiauth port mode</b> <i>port-string</i>
7.	Reset the ports MultiAuth authentication port maximum number of users to the default value for the specified ports.	<b>clear multiauth port numusers</b> <i>port-string</i>

### Setting MultiAuth Authentication Timers

The idle timeout setting determines the amount of idle time in which no traffic transits the link for a user or device before the connection is removed from the connection table. The idle timeout can be set for any authentication method.

The session timeout setting determines the maximum amount of time a session can last before being terminated.

[Procedure 43-12](#) describes setting the MultiAuth authentication timers.

### Procedure 43-12 MultiAuth Authentication Timers Configuration

Step	Task	Command(s)
1.	Optionally set the MultiAuth authentication idle timeout value for the specified authentication method.	<b>set multiauth idle-timeout</b> <i>auth-method timeout</i>
2.	Reset the MultiAuth authentication idle timeout value to its default value for the specified authentication method.	<b>clear multiauth idle-timeout</b> <i>auth-method</i>
3.	Optionally set the maximum amount of time a session can last before termination for the specified authentication method.	<b>set multiauth session-timeout</b> <i>auth-method timeout</i>
4.	Reset the maximum amount of time a session can last before termination to the default value for the specified authentication method.	<b>clear multiauth session-timeout</b> <i>auth-method</i>

## Setting MultiAuth Authentication Traps

Traps can be enabled at the system and module levels when the maximum number of users for the system and module, respectively, have been reached. Traps can be enabled at the port level for authentication success, failure, termination and when the maximum number of users have been reached on the port or all supported traps.

[Procedure 43-13](#) describes setting the MultiAuth authentication traps.

### Procedure 43-13 MultiAuth Authentication Traps Configuration

Step	Task	Command(s)
1.	Optionally enable MultiAuth authentication system traps.	<b>set multiauth trap system</b> { <b>enabled</b>   <b>disabled</b> }
2.	Optionally enable MultiAuth authentication module traps.	<b>set multiauth trap module</b> { <b>enabled</b>   <b>disabled</b> }
3.	Optionally enable MultiAuth authentication port traps.	<b>set multiauth trap port</b> <i>port-string</i> { <b>all</b>   <b>success</b>   <b>failed</b>   <b>terminated</b>   <b>max-reached</b> }
4.	Disable MultiAuth authentication traps for the specified trap type.	<b>clear multiauth trap</b> { <b>system</b>   <b>module</b>   <b>port</b> <i>portstring</i> { <b>all</b>   <b>success</b>   <b>failed</b>   <b>terminated</b>   <b>max-reached</b> }}

## Setting the MultiAuth Re-Authentication Timeout Action

When attempting to re-authenticate and the RADIUS server is unavailable, the re-authentication attempt times out and by default the session terminates. The re-authentication timeout action can be set to none. If the re-authentication timeout action is set to none, the session remains authenticated and provisioned according to the prior successful RADIUS authentication response.

It is recommended that you not set the re-authentication timeout action to none when using 802.1x authentication.

Use the **set multiauth reauth-timeout-action {terminate | none}** command to set the re-authentication timeout action for the device.

## Displaying MultiAuth Configuration Information

MultiAuth authentication supports the display of system-wide MultiAuth authentication values, MultiAuth authentication counters, port settings, end-user MAC addresses, session information, idle timeout settings, session timeout settings, and trap settings.

Table 43-2 describes displaying of MultiAuth authentication settings and statistics.

**Table 43-2 MultiAuth Authentication Settings and Statistics Display**

Task	Command(s)
Display system-wide MultiAuth authentication values.	<b>show multiauth</b>
Display MultiAuth authentication counters.	<b>show multiauth counters</b> [[ <b>quarantine-agent</b>   <b>cep</b>   <b>dot1x</b>   <b>mac</b>   <b>pwa</b>   <b>radius-snooping</b>   <b>auto-tracking</b> ] [ <b>chassis</b> ]   <b>port</b> <i>port-string</i> ]]
Display MultiAuth authentication port settings for all or the specified ports.	<b>show multiauth port</b> [ <i>port-string</i> ]
Display end-user MAC addresses per port for all MAC addresses and ports or for those specified.	<b>show multiauth station</b> [ <i>mac-address</i> ] [ <i>port-string</i> ]
Display MultiAuth authentication sessions for all sessions or the specified authentication method, MAC address, or ports.	<b>show multiauth session</b> [ <b>all</b> ] [ <b>agent</b> { <b>quarantine-agent</b>   <b>dot1x</b>   <b>mac</b>   <b>pwa</b>   <b>cep</b>   <b>radius-snooping</b>   <b>auto-tracking</b> }] [ <b>mac</b> <i>address</i> ] [ <b>port</b> <i>port-string</i> ]
Display MultiAuth authentication idle timeout values.	<b>show multiauth idle-timeout</b>
Display MultiAuth authentication session timeout values.	<b>show multiauth session-timeout</b>
Display MultiAuth authentication trap settings.	<b>show multiauth trap</b>

## Configuring VLAN Authorization

VLAN authorization allows for the dynamic assignment of users to the same VLAN. You configure VLAN authorization attributes within RADIUS. On the switch you enable VLAN authorization both globally and per-port. VLAN authorization is disabled globally by default. VLAN authorization is enabled per port by default. You can also set the VLAN egress format per-port. VLAN egress format defaults to un-tagged.

VLAN egress format can be set as follows:

- **none** – No egress manipulation will be made.
- **tagged** – The authenticating port will be added to the current tagged egress for the VLAN-ID returned.
- **untagged** – The authenticating port will be added to the current untagged egress for the VLAN-ID returned.
- **dynamic** – Egress formatting will be based upon information contained in the authentication response.

The VLAN authorization table will always list any tunnel attribute's VID's that have been received for authenticated end systems, but a VID will not actually be assigned unless VLAN authorization is enabled both globally and on the authenticating port. Dynamic VLAN authorization overrides the port PVID. Dynamic VLAN authorization is not reflected in the **show port vlan** display. The

VLAN egress list may be statically configured, enabled based upon the **set vlanauthorization egress** command, or have dynamic egress enabled to allow full VLAN membership and connectivity.

[Procedure 43-14](#) describes setting VLAN authorization configuration.

#### Procedure 43-14 VLAN Authorization Configuration

Step	Task	Command(s)
1.	Enable or disable VLAN authorization both globally and per port.	<b>set vlanauthorization {enable   disable}</b>
2.	Reset VLAN authorization configuration to default values for the specified port-list or for all.	<b>clear valanauthorization {port-list   all}</b>
3.	Display VLAN authorization configuration settings for the specified port-list or for all.	<b>show vlanauthorization {port-list   all}</b>

#### Setting Dynamic Policy Profile Assignment and Invalid Policy Action

Dynamic policy profile assignment is implemented using the policy mapping table. When VLAN authorization is enabled, authenticated users are dynamically assigned to the received tunnel attribute's VID, unless preempted by a policy map-table configuration entry. Dynamic policy profile assignment is supported by mapping a VID to a policy role upon receipt of a RADIUS tunnel attribute.

If the authentication server returns an invalid policy or VLAN to a switch for an authenticating supplicant, an invalid action of forward, drop, or default policy can be configured.

[Procedure 43-15](#) describes setting dynamic policy profile assignment and invalid policy action configuration.

#### Procedure 43-15 Policy Profile Assignment and Invalid Action Configuration

Step	Task	Command(s)
1.	Identify the profile index to be used in the VID-to-policy mapping.	<b>show policy profile all</b>
2.	Map the VLAN ID to the profile index.	<b>set policy mactable {vlan-list profile-index   response {tunnel   policy   both}}</b>
3.	Display the current mactable configuration.	<b>show policy mactable.</b>
4.	Set the action to take when an invalid policy or VLAN is received by the authenticating switch.	<b>set policy invalid action {default-policy   drop   forward}</b>

## Configuring RADIUS

You can set, clear, and display RADIUS configuration for both authentication and accounting.

### Configuring the Authentication Server

There are four aspects to configuring the authentication server:

- **State** enables or disables the RADIUS client for this switch.
- **Establishment values** configure a timer setting the length of time before retries, as well as the number of retries, before the switch determines the authentication server is down and attempts to establish with the next server in its list.

- **Server identification** provides for the configuration of the server IP address and index value. The index determines the order in which the switch will attempt to establish a session with an authentication server. After setting the index and IP address you are prompted to enter a secret value for this authentication server. Any authentication requests to this authentication server must present the correct secret value to gain authentication.
- The **realm** provides for configuration scope for this server: management access, network access, or both.

The 7100-Series firmware supports the configuration of multiple ASs. The lowest index value associated with the server determines the primary server. If the primary server is down, the operational server with the next lowest index value is used. If the switch fails to establish contact with the authentication server before a configured timeout, the switch will retry for the configured number of times.

Servers can be restricted to management access or network access authentication by configuring the realm option.

[Procedure 43-16](#) describes authentication server configuration.

### Procedure 43-16 Authentication Server Configuration

Step	Task	Command(s)
1.	Configure the index value, IP address, and secret value for this authentication server.	<b>set radius server</b> <i>index ip-address</i> [ <i>secret-value</i> ]
2.	Optionally set maximum number of sticky round robin authentication sessions allowed for either the specified RADIUS server or all RADIUS servers	<b>set radius max-sessions max-sessions</b> { <i>index</i>   <b>all</b> }
3.	Optionally set the number of seconds the switch will wait before retrying authentication server establishment.	<b>set radius timeout</b> <i>timeout</i>
4.	Optionally set the number of retries that will occur before the switch declares an authentication server down.	<b>set radius retries</b> <i>retries</i>
5.	Optionally set the authentication server configuration scope to management access, network access, or both for all or the specified authentication server.	<b>set radius realm</b> { <b>management-access</b>   <b>network-access</b>   <b>any</b> } { <i>as-index</i>   <b>all</b> }
6.	Optionally set the RADIUS authentication algorithm method for RADIUS server selection.	<b>set radius algorithm</b> { <b>standard</b>   <b>round-robin</b>   <b>sticky-round-robin</b> }
7.	Globally enable or disable RADIUS on the switch.	<b>set radius</b> { <b>enable</b>   <b>disable</b> }
8.	Reset the specified RADIUS setting to its default value.	<b>clear radius</b> {[ <i>state</i> ] [ <i>retries</i> ] [ <i>timeout</i> ] [ <i>server</i> [ <i>index</i>   <b>all</b> ] [ <i>realm</i> { <i>index</i>   <b>all</b> }]}
9.	Display the current RADIUS authentication server settings.	<b>show radius</b> [ <b>state</b>   <b>retries</b>   <b>authtype</b>   <b>timeout</b>   <b>server</b> [ <i>index</i>   <b>all</b> ]]

## Configuring RADIUS Accounting

There are four aspects to configuring RADIUS accounting:

- **State** enables or disables RADIUS accounting
- **Update values** allow the specification of the length of the period before accounting updates start and the interval between updates



- **Establishment values** configure a timer setting the length of time before retries, as well as the number of retries, before the switch determines the RADIUS accounting server is down and attempts to establish with the next server in its list.
- **Server identification** provides for the configuration of the RADIUS accounting server IP address and index value. The index determines the order in which the switch will attempt to establish with an accounting server. After setting the index and IP address you are prompted to enter a secret value for this accounting server.

Firmware supports the configuration of multiple RADIUS accounting servers. The lowest index value associated with the server determines the primary server. If the primary server is down, the operational server with the next lowest index value is used. If the switch fails to establish contact with the primary server before a configured timeout, the switch will retry for the configured number of times.

[Procedure 43-17](#) describes RADIUS accounting configuration.

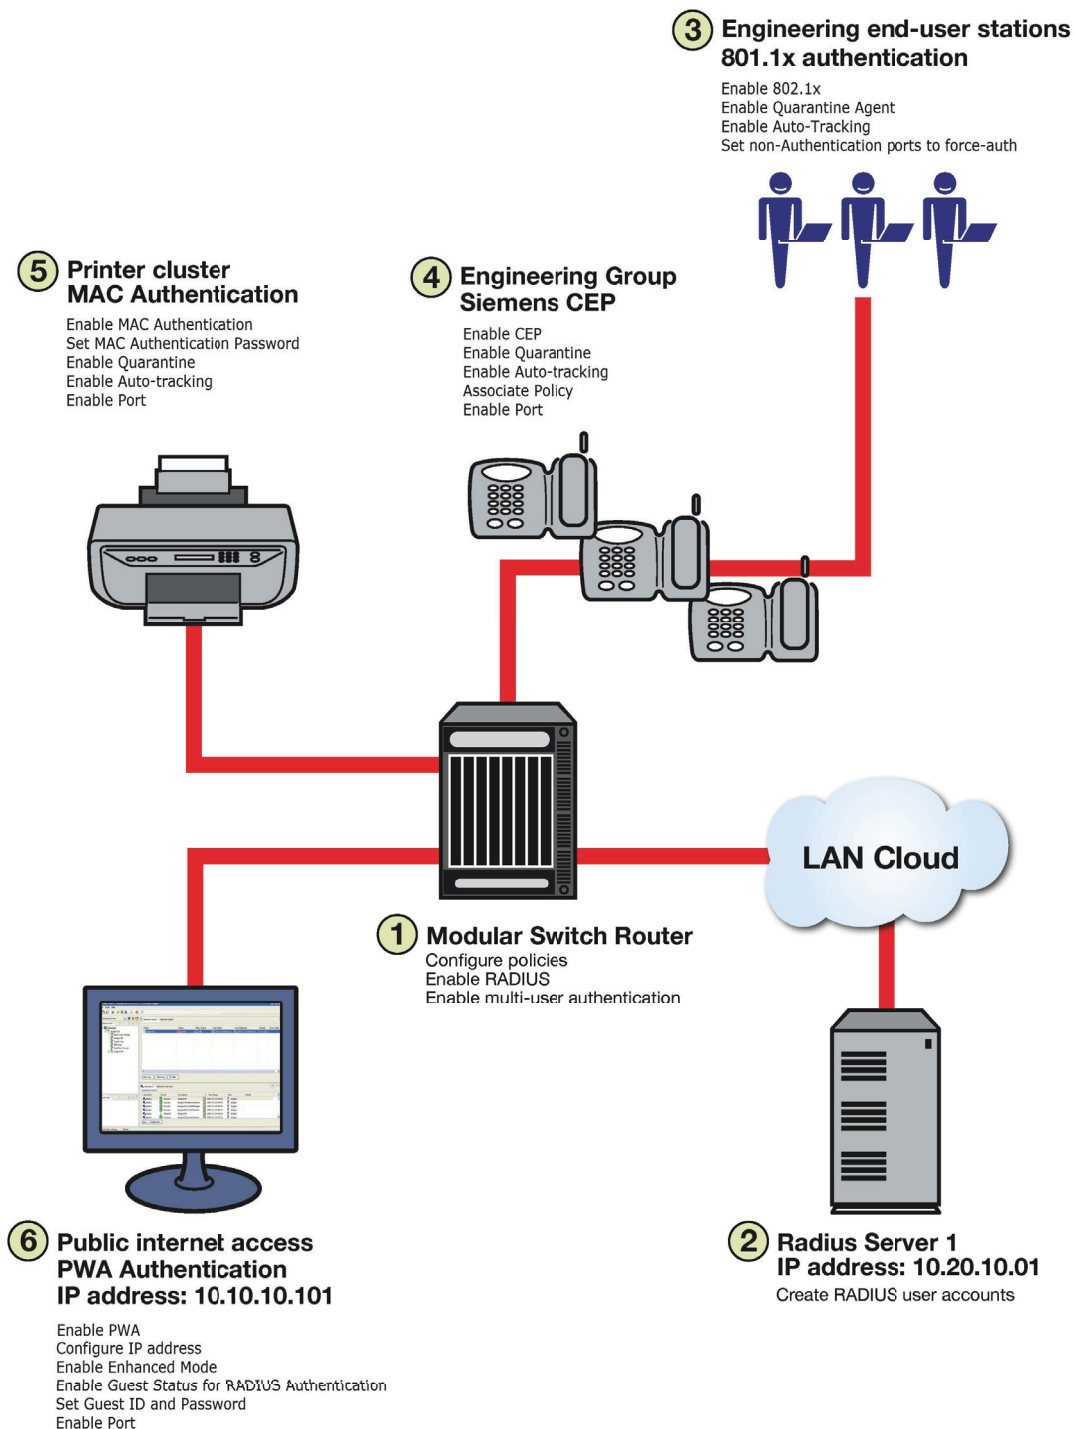
### Procedure 43-17 RADIUS Accounting Configuration

Step	Task	Command(s)
1.	Set the minimum interval at which RADIUS accounting sends interim updates.	<b>set radius accounting interval</b> <i>minimum interval</i>
2.	Set the number of seconds between each RADIUS accounting interim update.	<b>set radius accounting update</b> <i>interval interval</i>
3.	Set the number of times a switch will attempt to contact a RADIUS accounting server.	<b>set radius accounting retries</b> <i>retries</i>
4.	Set the amount of time to establish contact with a RADIUS accounting server before timing out.	<b>set radius accounting timeout</b> <i>timeout {index   all}</i>
5.	Configure the RADIUS accounting server.	<b>set radius accounting server</b> <i>{index   all} ip_address udp-port [server-secret]</i>
6.	Enable or disable RADIUS accounting on this switch.	<b>set radius accounting</b> <i>{enable   disable}</i>
7.	Reset RADIUS accounting parameters to default values or clear server definitions on this switch.	<b>clear radius accounting</b> <i>{[server {index   all}] [retries {index   all}] [timeout {index   all}] [intervalminimum] [updateinterval]}</i>
8.	Display RADIUS accounting configuration or statistics.	<b>show radius accounting</b> <i>[updateinterval   intervalminimum   state   server {index   all}]</i>

# Authentication Configuration Example

Our example covers six supported authentication types being used in an engineering group context: an end-user station, an IP phone, a printer cluster, and public internet access, along with both the quarantine and auto-tracking agents turned on in appropriate places.

Figure 43-4 Authentication Configuration Example Overview



Our configuration example consists of the following steps as shown in [Figure 43-4](#) and described in the sections that follow:

1. Configuring policies, RADIUS, and MultiAuth authentication on the switch.
2. Creating RADIUS user accounts on the authentication server.
3. Configuring for the engineering group 802.1x end-user stations.
4. Configuring for the engineering group Siemens CEP devices.
5. Configuring for the printer cluster MAC authentication.
6. Configuring for the public area internet access for PWA.

## Configuring the Quarantine Agent

We will enable the quarantine agent on the switch, but only turn it on for the 802.1x, MAC, and CEP authentication contexts.

For a quarantine policy example we will create a rule for forwarding UDP source port 67 which is normally used for DHCP traffic and associate it with the dhcpQuarantine policy profile. We want to disable any port not connected to a DHCP server if attempts to forward any DHCP traffic occurs. The following CLI input

- Enables the quarantine agent on the switch
- Sets policy rule 1 to hit when forwarding traffic occurs on UDP source port 67 and specifies that the rule is assigned to policy profile 1 as a quarantine profile.
- Names the policy profile dhcpQuarantine and sets the action to disable the port should a hit occur for rule 1.

We enable the quarantine agent at the port level within the appropriate 802.1, MAC, and CEP authentication discussions.

```
7100 Chassis(rw)->set quarantine-agent enable
7100 Chassis(rw)->set policy rule 1 udpportsourceip 67 mask 16 forward
quarantine-profile 1
7100 Chassis(rw)->set policy profile 1 name dhcpQuarantine disable-port enable
```

## Configuring the Auto-Tracking Agent

We will enable the auto-tracking agent on the switch, but only turn it on for the 802.1x, MAC, and CEP authentication contexts. We want to enable enhanced mode for PWA. PWA enhanced mode is not supported if auto-tracking is enabled on the port. If you configure an admin-policy, it will be used should auto-tracking be selected as the authenticator. If you do not configure an admin-policy, the device will authenticate, the session will be logged, and no action will occur.

The following CLI input enables the auto-tracking agent on the switch. We enable the auto-tracking agent at the port level within the appropriate 802.1, MAC, and CEP authentication discussions.

```
7100 Chassis(rw)->set auto-tracking enable
```

## Setting MultiAuth Configuration On the Switch

MultiAuth authentication must be set to **multi** whenever multiple users of 802.1x need to be authenticated or whenever any non-802.1X authentication method is present. For ports where no authentication is present, such as switch to switch, or switch to router connections, you should also set MultiAuth port mode to force authenticate to assure that traffic is not blocked by a failed

authentication. For purposes of this example, we will limit authentication to a maximum of 6 users per port.

The following CLI input:

- Sets MultiAuth authentication to **multi**.
- Sets ports with switch to switch and switch to router connections to force authenticate.
- Sets the maximum number of users that can authenticate on each port to 6.

```
7100 Chassis(rw)->set multiauth mode multi
7100 Chassis(rw)->set multiauth port mode force-auth ge.1.5-7
7100 Chassis(rw)->set multiauth port numusers 6 ge.1.5-7
7100 Chassis(rw)->set multiauth port mode force-auth ge.1.19-24
7100 Chassis(rw)->set multiauth port numusers 6 ge.1.19-24
```

- Enables MultiAuth authentication system and module traps for the 7100-Series configuration.

```
7100 Chassis(rw)->set multiauth trap system enabled
7100 Chassis(rw)->set multiauth trap module enabled
```

This completes the MultiAuth authentication configuration piece for this example. Keep in mind that you would want to use the **set multiauth precedence** command, to specify which authentication method should take precedence, should you have a single user configured for multiple authentications on the same port.

## Enabling RADIUS On the Switch

The switch needs to be informed about the authentication server. Use the following CLI input to

- Configure the authentication server IP address on the switch.
- Enable the RADIUS server.

```
7100 Chassis(rw)->set radius server 1 10.20.10.01
7100 Chassis(rw)->set radius enable
```

## Creating RADIUS User Accounts On The Authentication Server

RADIUS account creation on the authentication server is specific to the RADIUS application you are using. Please see the documentation that comes with your RADIUS application. Create an account for all users to be authenticated.

## Configuring the Engineering Group 802.1x End-User Stations

There are three aspects to configuring 802.1x for the engineering group:

- Configure EAP on each end-user station
- Set up an account in RADIUS on the authentication server for each end-user station
- Configure the quarantine agent, 802.1x, and the auto-tracking agent on the switch

Configuring EAP on the end-user station and setting up the RADIUS account for each station is dependent upon your operating system and the RADIUS application being used, respectively. The important thing the network administrator should keep in mind is that these two configurations should be in place before moving on to the 802.1x configuration on the switch. In an 802.1x configuration, policy is specified in the RADIUS account configuration on the authentication server using the RADIUS Filter-ID. See [“The RADIUS Filter-ID”](#) on page 43-10 for RADIUS

Filter-ID information. If a RADIUS Filter-ID exists for the user account, the RADIUS protocol returns it in the RADIUS Accept message and the firmware applies the policy to the user.



**Note:** Globally enabling 802.1x on a switch sets the port-control type to **auto** for all ports. Be sure to set port-control to **forced-auth** on all ports that will not be authenticating using 802.1x and no other authentication method is configured. Otherwise these ports will fail authentication and traffic will be blocked.

The following CLI input:

- Enables 802.1x on the switch
- Sets port-control to **forced-auth** for all connections between switches and routers, because they do not use authentication and would be blocked if not set to **forced-auth**
- Enables the quarantine agent on ports ge.1.5, ge.1.19, and ge.1.24
- Enables the auto-tracking agent on the switch and ports ge.1.5, ge.1.19, and ge.1.24

```
7100 Chassis(rw)->set dot1x enable
7100 Chassis(rw)->set dot1x auth-config authcontrolled-portcontrol forced-auth
ge.1.5
7100 Chassis(rw)->set dot1x auth-config authcontrolled-portcontrol forced-auth
ge.1.19
7100 Chassis(rw)->set dot1x auth-config authcontrolled-portcontrol forced-auth
ge.2.24
7100 Chassis(rw)->set quarantine-agent port enable ge.1.5
7100 Chassis(rw)->set quarantine-agent port enable ge.1.19
7100 Chassis(rw)->set quarantine-agent port enable ge.1.24
7100 Chassis(rw)->set auto-tracking enable
7100 Chassis(rw)->set auto-tracking port enable ge.1.5
7100 Chassis(rw)->set auto-tracking port enable ge.1.19
7100 Chassis(rw)->set auto-tracking port enable ge.1.24
```

This completes the 802.1x end-user stations configuration.

## Configuring the Engineering Group Siemens CEP Devices

If a Siemens phone is inserted into a port enabled for Siemens CEP, the firmware detects communication on UDP/TCP port 4060. Use policy manager to configure a policy with a VLAN, CoS, and rate limit appropriate to VoIP. See [Chapter 40, Quality of Service \(QoS\) Configuration](#) for a QoS VoIP policy configuration example. Once an existing policy is configured, the **set cep policy command can be used to apply** the policy

The following CLI input:

- Enables CEP globally on the switch.
- Sets CEP policy to a previously configured policy named **siemens** with an index of **9**.
- Sets ports **ge.1.16-18** to only accept default Siemens type phones and applies the Siemens policy to the specified ports.
- Enable the quarantine and auto-tracking agents on ports **ge.1.16** through **ge.1.18**

```
7100 Chassis(rw)->set cep enable
7100 Chassis(rw)->set cep policy siemens 9
7100 Chassis(rw)->set cep port ge.1.16-18 siemens enable
7100 Chassis(rw)->set quarantine-agent port enable ge.1.16-18
```

```
7100 Chassis(rw)->set auto-tracking port enable ge.1.16-18
```

This completes the Siemens CEP end-user stations configuration.

## Configuring the Printer Cluster for MAC-Based Authentication

Perform the following tasks to configure MAC-based authentication for the printer cluster in our example:

- Set up an account for each printer on the authentication server that contains the printer MAC address, the MAC authentication password configured on the switch, and a RADIUS filter ID entry specifying the printer policy.
- Configure a policy using the policy manager specifying the printer cluster VLAN and optionally configuring a CoS and rate limit.
- Enable MAC authentication globally on the switch.
- Enter the MAC authentication password as **enterasys** on the switch.
- Set the MAC authentication significant-bits to **24**.
- Enable MAC authentication on the ports used by the printer cluster: **ge.1.3-4**
- Enable the quarantine and auto-tracking agents on ports **ge.1.3-4**

With the authentication server configured with a RADIUS account for each printer, and the printer policy preconfigured, enter the following CLI input:

```
7100 Chassis(rw)->set macauthentication enable
7100 Chassis(rw)->set macauthentication password enterasys
7100 Chassis(rw)->set macauthentication significant-bits 24
7100 Chassis(rw)->set macauthentication port enable ge.1.3-4
7100 Chassis(rw)->set quarantine-agent port enable ge.1.3-4
7100 Chassis(rw)->set auto-tracking port enable ge.1.3-4
```

This completes the printer cluster MAC authentication configuration.

## Configuring the Public Area PWA Station

The public area PWA station provides visitors to your business site with open access to the internet, while at the same time isolating the station from any access to your internal network. In order to provide a default set of network resources to communicate over HTTP, policy must be set to only allow DHCP, ARP, DNS, and HTTP. You may want to set a rate limit that would guard against excessive streaming. You will also need to set up RADIUS for the public station account on the authentication server. This configuration will include the guest name, password, and a RADIUS Filter-ID for the public policy. We will not enable auto-tracking because PWA enhanced mode is not supported with auto-tracking. We will also not enable quarantine.

Perform the following tasks to configure the public station for PWA authentication:

- Configure the policy appropriate to the public station.
- Setup the RADIUS user account for the public station on the authentication server.
- Enable PWA globally on the switch.
- Configure the IP address for the public station.
- Optionally set up a banner for the initial PWA screen.
- Enable PWA enhancemode so that any URL input will cause the PWA sign in screen to appear.

- Set PWA gueststatus to RADIUS authentication mode.
- Set the PWA login guest name.
- Set the PWA login password.
- Enable PWA on the switch port where the public station is connected.

Once the policy and RADIUS account are configured, enter the following CLI input on the switch:

```
7100 Chassis(rw) ->set pwa enable
7100 Chassis(rw) ->set pwa ipaddress 10.10.10.101
7100 Chassis(rw) ->set pwa banner \"Extreme Networks Public Internet Access
Station\"
7100 Chassis(rw) ->set pwa enhancemode enable
7100 Chassis(rw) ->set pwa gueststatus authradius
7100 Chassis(rw) ->set pwa guestname guest
7100 Chassis(rw) ->set pwa guestpassword password
7100 Chassis(rw) ->set pwa portcontrol enable ge.1.6
```

This completes the Authentication configuration example.

## Terms and Definitions

[Table 43-3](#) lists terms and definitions used in this Authentication configuration discussion.

**Table 43-3 Quality of Service Configuration Terms and Definitions**

Term	Definition
Authentication Server (AS)	An entity providing authorization services to an authenticator using RADIUS. The authentication server may be on the same device or be at a remote location.
Authenticator	The switch seeking authentication from the authentication server for a supplicant.
Auto-tracking agent	A form of authentication that authenticates those sessions that are not captured by the other supported MultiAuth authentication agents (quarantine, 802.1x, PWA, MAC, CEP, and RADIUS snooping).
Convergence End Point (CEP)	A protocol capable of detecting an IP telephony or video device on a port and dynamically applying a specific policy to the port.
Domain Name System (DNS)	Serves as a means for the Internet to translate human-readable computer hostnames, e.g. www.example.com, into the IP addresses.
Dynamic Host Configuration Protocol (DHCP)	A protocol used by networked clients to obtain various parameters necessary for the clients to operate in an Internet Protocol (IP) network.
Extensible Authentication Protocol (EAP)	A protocol that provides the means for communicating the authentication information in an IEEE 802.1x context.
IEEE 802.1x	An IEEE standard for port-based Network Access Control that provides authentication to devices attached to a LAN port, establishing a point-to-point connection or preventing access from that port if authentication fails.
MAC-based Authentication	A means of authenticating a device attempting to gain access to the network based upon the device MAC address and a secret keyword known to the authenticator and the RADIUS application on the authentication server.
Multi-user Authentication	The ability to appropriately authenticate multiple supplicants on a single link and provision network resources, based upon policy associated with each supplicant.

**Table 43-3 Quality of Service Configuration Terms and Definitions (continued)**

<b>Term</b>	<b>Definition</b>
MultiAuth Authentication	The ability to authenticate multiple authentication modes for a user and applying the authentication mode with the highest precedence.
Port Web Authentication (PWA)	A means of authenticating a user by utilizing a web browser for the login process to authenticate to the network.
Quarantine agent	A form of authentication that depends upon the existence of one or more configured quarantine policy rules, with each rule associated with a policy profile that determine the action should the quarantine agent be used to authenticate the device.
RADIUS Filter ID	An Extreme Networks proprietary string formatted in the RADIUS Access-Accept packet sent back from the authentication server to the switch containing either the policy to apply to the supplicant, the management type for the port, or both.
RADIUS Protocol	An AAA (Authentication, Authorization, and Accounting) protocol for controlling access to network resources used by ISPs and corporations managing access to Internet or internal networks across an array of access technologies.
Supplicant	The user or device seeking access to network resources.